Dell Compellent HP-UX Best Practices 11i v2 and 11i v1

Best Practices, Tips and Tricks

© 2012 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

Trademarks used in this text: Dell™, the DELL™ logo, and Compellent™ are trademarks of Dell Inc.

Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

May 2011 Rev. A

Contents

General syntax 5
Document Revision 5
Icons 6
Preface 7
Audience and Assumptions 7
Purpose 7
Important Notes 7
Customer support 7
Introduction 8
Best Practices 9
Setup and Configuration 11
HP-UX 11i v2 and v1 with Controller Failover 11
HP-UX 11i v2 and v1 Server Path Failover 11
Server Setup 11
LUN Discovery 12
Creating a new Volume Group 14
Creating a new Logical Volume 18
Mounting and using the new Logical Volume/Filesystem 19
Dynamic LUN Expansion (DLE) and Dynamic LUN Contraction (DLC) 21
Expanding the LUN 21
Exposing the new capacity to the Volume Group 22
Dynamic LUN Contraction (DLC) and Thoughts 26
Data Instant Replay (Continuous Snapshots) 27
Quiesce the Volume Group 28
Capturing the Snapshot 29
Mounting to the same host 33
Mounting to a alternate host 35
Working with Striped/Mirror or MultiVolume Volume Groups 36
Thin Provisioning 45
Tuning for Performance 46
Kernel Tuning 46
SCSI Configuration 46
Filesystem Configuration 46

Dynamic Root Disk 47
DRD and creating a Point in Time Copy 48
DRD and creating a Boot from SAN LUN 50
Boot from SAN (OS Boot Volume) 55
iSCSI and HP-UX 56
Assumptions 56
Setup 56
Adding LUNs 59
CompCU (Dell Compellent Command Utility) 62
Setup and Validation 62
Example 1: Creating a Single Volume 64
Example 2: Creating Multiple Volumes 65
Example 3: Capturing a Volume Replay 66
FAQ 67

General syntax

Table 1. Conventions

Item	Convention
Menu items, dialog box titles, field names, keys	**Bold**
Mouse click required	Click
User Input	`Monospace Font`
User typing required	Type:
System response to commands	Blue
Website addresses	http://www.dell.com
Email addresses	name@dell.com

Document Revision

Table 2. Revision History

Date	Revision	Comments
04/02/2012		Updates incl. LUN 0 requirements, Boot and Install to SAN volumes
10/14/2011	A	Initial Draft

Icons

Notes are used to convey special information or instructions.

Timesavers are tips specifically designed to save time or reduce the number of steps.

Caution indicates the potential for risk including system or data damage.

Warning indicates that failure to follow directions could result in bodily harm.

Preface

Audience and Assumptions

- The audience is assumed to be junior to mid/level system administrators who are not necessarily familiar with HP-UX. The user is:
  - Already trained on the usage of the **Dell Compellent Storage** or storage has already been allocated; this is not covered here.
  - Familiar with switch zoning or the zoning is already accomplished as this is not covered here.
  - Familiar with the documents referenced in the Introduction section below.
  - Registered with and has access to the HP Support Center (formerly HP ITRC, located here http://h20566.www2.hp.com/portal/site/hpsc/public/). Registration is free, but access to software updates and the patch database would require the registration of a properly licensed and valid Support Contract ID.
- Firmware and/or EFI levels on all hardware (servers/HBAs/etc.) are within six months of current levels.
- HP-UX 11i environment is patched to a reasonable level (i.e., QPK patch sets are not more than six months old OR last known best patch set are installed in the case of EOL software (i.e. HP-UX 11iv1, Dec 2006).
- All statements contained herein are based upon a test platform which consists of one (1) HP-UX Itanium rx3600 server (with 4 processors, 16GB RAM, 4x 1000bT network ports, AD193A FC HBA and AD194A FC HBA), built with HP-UX 11i v2 and Dec 2010 Support Plus installed.

Purpose

This document provides an overview and outlines the recommended Best Practices, Tips and Tricks, as advised by Dell Compellent when using its Storage Center products (v5.4.x, v5.5.x and higher) in and with HP-UX 11i environments (both ia64 and PA-RISC architectures).

Important Notes

- With Dell Compellent Storage Center v5.3.x and older and HP-UX 11.00, 11i v1 and 11i v2, device addressing is limited to 8 LUNs per HBA, with LUN IDs numbered 0 to 7. There must be a LUN 0 **and** you cannot skip LUN 0 under any circumstance!
- HP-UX 11.00, HP-UX 11i v1 and HP-UX 11i v2 do not have native multi-path Input/Output (MPIO) capabilities; redundant link failover **is** built in to the native Logical Volume Manager and is called "PVLinks". This is discussed both here and in other Dell Compellent documents.
- HP-UX 11i v3 has native MPIO capabilities; referred to as the "Agile" device naming system.
- Please check for and review any Copilot Service Technical Alerts (CSTA) for your specific Operating System at Dell Compellent's Knowledge Center **BEFORE** attempting any configuration additions or changes.

Customer support

Dell Compellent provides live support 1-866-EZSTORE (866.397.8673), 24 hours a day, 7 days a week, 365 days a year. For additional support, email Dell Compellent at support@compellent.com. Dell Compellent responds to emails during normal business hours.

Introduction

This document presents useful information on a wide variety of HP-UX topics not covered in other Dell Compellent documentation. While this document contains significant detail, it is by no means a comprehensive or complete representation of every HP-UX method or function (nor is it intended to be). All of the procedures contained herein are handled more completely in HP documentation, which also includes the system MAN pages. It is strongly recommended that you obtain the following HP documents and familiarize yourself with basic system commands and functions.

- Core HP-UX, Manpages http://www.hp.com/go/hpux-core-docs
- I/O Cards http://www.hp.com/go/hpux-iocards-docs
- I/O Cards for Integrity http://www.hp.com/go/integrity-iocards-docs
- Serviceguard Solutions http://www.hp.com/go/hpux-serviceguard-docs
- Virtualization http://www.hp.com/go/virtualization-manuals

With HP-UX 11iv3 (i.e. 11.31), there have been significant changes made to the mass storage stack. It is recommended that you obtain the following documents from HP and familiarize yourself with all of the new changes incorporated in this version of HP-UX, located here:

- LVM New Features, 11i v3
  http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01919369/c01919369.pdf
- The Next Generation Mass Storage Stack
  http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01915888/c01915888.pdf
- 11i v2 to 11i v3 Mass Storage Stack Update Guide
  http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01915888/c01915888.pdf
- LVM Migration from Legacy to Agile Naming Model, HP-UX 11i v3
  http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01916036/c01916036.pdf
- HP-UX 11i v3 Mass Storage Device Naming
  http://bizsupport.austin.hp.com/bc/docs/support/SupportManual/c01906504/c01906504.pdf
- HP-UX 11i v3 Native Multi-Pathing for Mass Storage
  http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01906512/c01906512.pdf
- SCSI Management and Diagnostics, HP-UX 11i v3
  http://bizsupport2.austin.hp.com/bc/docs/support/SupportManual/c01914672/c01914672.pdf

Best Practices

Best practices in enterprise environments dictate the elimination of single points of failure in server installations; this includes all required resources: power, storage and networking. For purposes of this document, we are assuming that both networking and power are redundant and configured correctly. We are solely concerned with the implementation of Dell Compellent storage in an enterprise environment with HP-UX 11i v3, v2 and/or v1 servers.

To that end, the optimum configuration for HP-UX servers accessing Dell Compellent storage consists of at least one pair of Storage Center heads, two separate Fibre Channel (FC) fabrics and at least two HBAs (host bus adapters) in each HP-UX server. There are major differences between the 'mass storage stack' on HP-UX 11i v3 and all earlier versions of HP-UX; for that reason, this document is written to address HP-UX 11i v2, v1 and 11.00 only, with the accompanying document "**Dell Compellent HP-UX Best Practices 11i v3**" written to address the differences and nuances of v3 accordingly.

Figure 1. HP-UX Best Practices Configuration

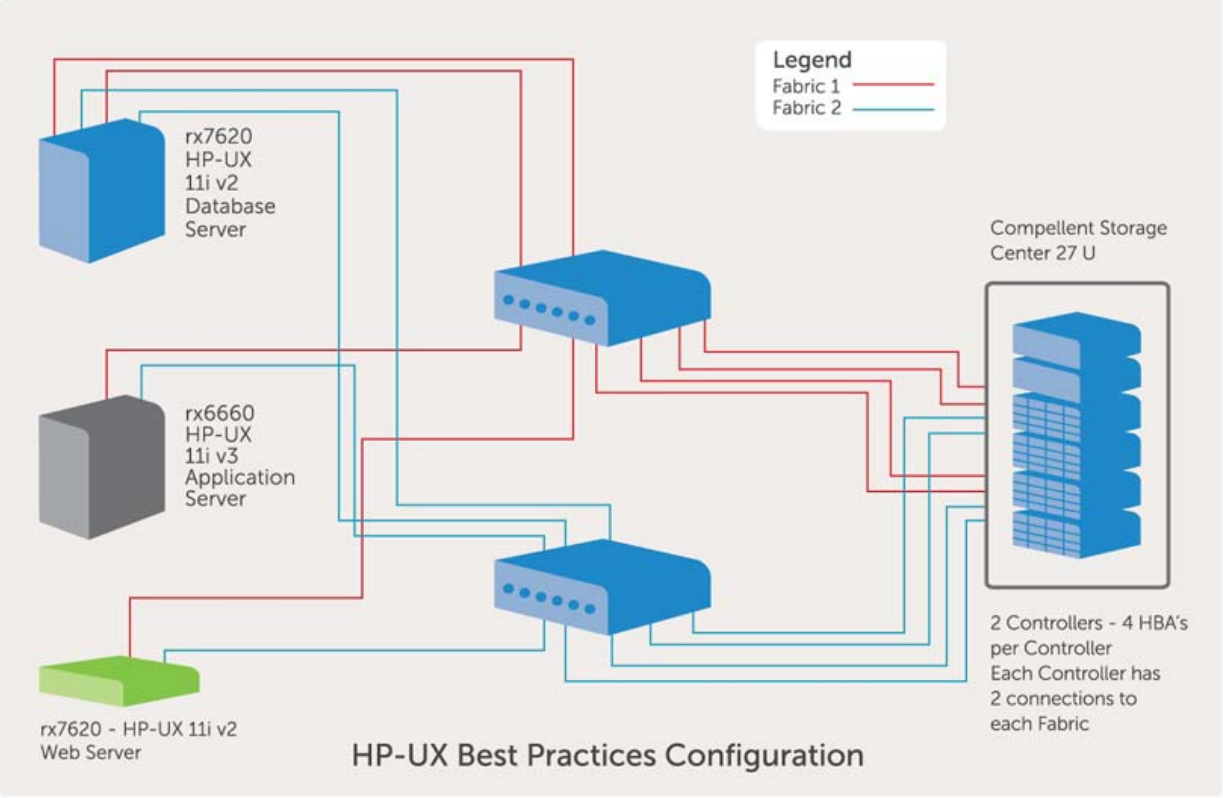

Note

The above scenario will continue to function even if the HP-UX server may contain ONLY one (1) FC port on a single FC HBA. In this latter scenario with limited hardware, all LUNs from the Dell Compellent Storage Center are still presented in a redundant manner across each Dell

Compellent Storage Center controller, LUNs are also mapped across the FC fabric in a redundant manner (i.e. one or more VLAN); although all mappings would tie into the single WWPN/WWNN of the single FC port on the HP-UX server. From the HP-UX server standpoint, all LUNs discovered via the OS will continue to be dual/pathed and any failure at the Dell Compellent Storage Center controller, the FC fabric or FC fabric port will be tolerated with the HP-UX server not suffering any IO impact or loss. This scenario will fall apart if either the single FC port or single FC HBA on the HP-UX does fail; and would be the single point of failure in this limited hardware configuration.

Setup and Configuration

HP-UX 11i v2 and v1 with Controller Failover

Dell Compellent Storage Center employs two methods for high availability.

- Legacy Mode
- N_Port ID Virtualization (NPIV)

For more in-depth information on these two modes, please refer to Dell Compellent Storage Center User Guide.

When attaching HP-UX 11i v2 and v1 servers to the Dell Compellent Storage Center, Legacy mode must be configured on the Storage Center; due to the device addressing scheme on HP-UX 11i v2 and v1 servers, NPIV mode on the Storage Center **SHOULD NOT** be used (also see Note below). In addition to the Legacy mode, the Fibre Channel switch that the Storage Center is plugged into must have port binding to WWN feature. For example, Cisco switches have Persistent FCID while Brocade switches have Persistent PID. Please refer to Dell Compellent CSTA for more details.

NPIV functionality is dependent upon the hardware and HBAs ability to understand the NPIV protocol. As of this paper, NPIV is supported on newer 8GB PCIe FC HBAs and HP BladeSystem c-Class blades ONLY.

When creating LUNs on the Dell Compellent Storage Center, the LUNs can only be mapped to a HP-UX server through one controller or the other but not through both at the same time. When this controller fails, the LUNs will automatically move to the second controller. This is the reason why port binding to WWN feature on the switch is required in order for the HP-UX server to maintain the same disk device files after the failover or failback.

HP-UX 11i v2 and v1 Server Path Failover

HP-UX 11i v2 and v1 uses PVLinks via Logical Volume Manager (LVM) as a way to handle server path failover. When creating LUNs on the Dell Compellent Storage Center, the LUNs must be mapped to both HBA ports on the target host.

Server Setup

In a dual/zone fabric, zone each host HBA via separate zones to the same four (4) Front End Primary (FEP) and four (4) Front End Reserve (FER) ports/domains.

If you need to determine the World Wide Port Names (WWPNs) of your host FC HBAs, then issue the following command.

```
Type: ls –la /dev/*d* | grep –e td –e fcd | awk ‘{print $10}’
/dev/fcd0
/dev/td1
Type: /opt/fcms/bin/fcmsutil /dev/td1 | grep Port
```

Local N_Port_id is = 0xe90600
N_Port Node World Wide Name = 0x50060b00002350c1
N_Port Port World Wide Name = 0x50060b00002350c0

The text highlighted in RED above, represents the WWPN of the /dev/td1 HBA; repeat this command with each HBA device path to determine their respective WWPN.

After the zoning is complete, run the following command on your server; this will register the HBA WWPNs into the Dell Compellent Storage Center table of known connections.

Type: `ioscan –fn`

On your Dell Compellent Storage Center, initiate a "Create Server" dialog, uncheck the "Only Show Up Connections" box and let the screen refresh. You should now be able to create your server object.

For a HP-UX 11i v2 or v1 server object, please also ensure that the Operating System type is set to "**HP UX 11i PVLinks**".

Create a 30GB LUN on the Dell Compellent Storage Center.

Prior to mapping the new LUN/s to the server object, login to the target host and capture the output of an ioscan command; for HP-UX 11i v2 or v1, run "**ioscan -fnC disk > /tmp/ioscan.b4"**.

Using the Dell Compellent Storage Center interface, map this new LUN to both host HBAs via the newly created and discovered HP-UX server object.

In this Advanced configuration dialog, you want to mark the checkbox which says "Create maps to down server ports", this checkbox is only available if the FC ports to the target host is in a Down state; then left-Click the "Continue" button.

LUN Discovery

Run the following command to force a rescan of all devices on the system.

Type: `ioscan –fn`

Now, run the following command to get a list of all "known" LUNs and disk devices.

Type: `ioscan –fnC disk > /tmp/ioscan.after`

`The difference between this file and the file previously captured (/tmp/ioscan.b4, see above Timesaver) would reflect any and all new LUNs`

presented from the Dell Compellent Storage Center to the host, as shown in RED below.

```
Class     I  H/W Path                    Driver   S/W State   H/W Type     Description
===================================================================================
disk      4  0/0/2/0.6.0                 sdisk    CLAIMED     DEVICE       SEAGATE ST318404LC
                                         /dev/dsk/c1t6d0       /dev/rdsk/c1t6d0
<snip>
disk      6  1/4/0/0.206.2.255.0.0.0     sdisk    CLAIMED     DEVICE       COMPELNTCompellent Vol
                                         /dev/dsk/c27t0d0      /dev/rdsk/c27t0d0
disk     28  1/4/0/0.206.3.255.0.0.0     sdisk    CLAIMED     DEVICE       COMPELNTCompellent Vol
                                         /dev/dsk/c28t0d0      /dev/rdsk/c28t0d0
disk     29  1/4/0/0.206.4.255.0.0.0     sdisk    CLAIMED     DEVICE       COMPELNTCompellent Vol
                                         /dev/dsk/c29t0d0      /dev/rdsk/c29t0d0
disk     30  1/4/0/0.206.5.255.0.0.0     sdisk    CLAIMED     DEVICE       COMPELNTCompellent Vol
                                         /dev/dsk/c30t0d0      /dev/rdsk/c30t0d0
```

Note

If no new devices are discovered or shown, then run the command "**insf -e**" to force a scan and rebuild of the system device files/device tree; thereafter repeating the ioscan command from above.

```
insf: Installing special files for sctl instance 2 address 0/0/1/0.7.0
insf: Installing special files for sdisk instance 4 address 0/0/2/0.6.0
<snip>
insf: Installing special files for sdisk instance 7 address 1/4/0/0.206.2.255.0.0.1
insf: Installing special files for sdisk instance 8 address 1/4/0/0.206.2.255.0.0.2
insf: Installing special files for pseudo driver diag2
insf: Installing special files for pseudo driver root
```

To validate that the new device (as discovered) is presented from the Dell Compellent Storage Center, issue the following command.

Type: `/usr/sbin/diskinfo /dev/rdsk/c27t0d0`

```
SCSI describe of /dev/rdsk/c27t0d0:
                 vendor:  COMPELNT
             product id:  Compellent Vol
                   type:  direct access
                   size:  31457280 Kbytes
       bytes per sector:  512
```

Timesaver

To confirm that this is the actual LUN, matched by its unique serial number, (where you

would correlate the “Hardware path:” of the LUN to its respective “Serial Number:”)

```
Type: echo "selclass type disk;info;wait;infolog" | /usr/sbin/cstm | egrep -e
"Hardware path" -e "Serial Number"
```

```
Hardware path: 0/0/2/1.0.16.0.0
Serial Number:    N/A
Hardware path: 0/2/1/0/4/0.236.0.0.0.0.1
Serial Number:    00000067-0000074c
Hardware path: 0/2/1/0/4/0.236.1.0.0.0.1
Serial Number:    00000067-0000074a
Hardware path: 0/2/1/0/4/0.236.1.0.0.0.2
Serial Number:    00000067-00000752
Hardware path: 0/2/1/0/4/1.236.7.0.0.0.1
Serial Number:    00000067-0000074c
Hardware path: 0/2/1/0/4/1.236.9.0.0.0.1
Serial Number:    00000067-0000074a
Hardware path: 0/2/1/0/4/1.236.9.0.0.0.2
Serial Number:    00000067-00000752
Hardware path: 1/4/0/0.206.2.255.0.0.0
Serial Number:    00000067-00000754
Hardware path: 0/4/1/0.0.0.0.0
Serial Number:        P4V5ZLYA
Hardware path: 0/4/1/0.0.0.1.0
Serial Number:        P4WPNABA
```

To use the new device, you will have to prepare it for LVM. Issue the following command.

```
Type: pvcreate /dev/rdsk/c27t0d0
```

```
Physical volume "/dev/rdsk/c27t0d0" has been successfully created.
```

Creating a new Volume Group

To create a new Volume Group (and manage the allocation of the Volume Group ID/number), use this procedure.

a. You need to create the directory structure where the Volume Group will exist and the group file used to identify it. In order to find out what Volume Groups already exist and are in use, issue the command.

```
Type: ls -la /dev/*/group
```

```
crw-r-----      1 root          sys             64  0x000000   Aug 13 03:10
        /dev/vg00/group
crw-rw-rw-      1 root          sys            128 0x001000   Dec  2 13:38
        /dev/vg01/group
```

b. It is good to note at this juncture the difference between the two device files. The major number 64 represents an LVM 1.x Volume Group (HP-UX 11i v2 or v1 Volume Groups), while the major number of 128 represents an LVM 2.x Volume Group (HP-UX 11i v3 Volume Groups). Also, with LVM 1.x, the first two minor numbers (as shown in RED above) represent the Volume Group number; with LVM 2.x, we use the first three. As an example, we will create a Volume Group called "vgdemo". First, create a directory with the same name in /dev with the following command.

```
Type: mkdir –p /dev/vgdemo
```

c. Next, create the device file you want to use to identify the Volume Group. As you see in RED above, both 00 and 001 are in use. We will thus choose 0x02000 (LVM 1.x) as the next Volume Group ID/number. To do this, issue the command.

```
Type: mknod /dev/vgdemo/group c 64 0x020000
Type: ls -la /dev/*/group

crw-r-----   1 root   sys      64 0x000000 Aug 13 03:10 /dev/vg00/group
crw-rw-rw-   1 root   sys     128 0x001000 Dec  2 13:38 /dev/vg01/group
crw-rw-rw-   1 root   sys      64 0x020000 Dec  3 14:43 /dev/vgdemo/group
```

d. Once complete, you can now issue the vgcreate command (as outlined above) to create the Volume Group.

```
Type: vgcreate –s 16 /dev/vgdemo /dev/dsk/c27t0d0
```

(where **-s**: use 16MB PE size)

```
Increased the number of physical extents per physical volume to 1919.
Volume group "/dev/vgdemo" has been successfully created.
Volume Group configuration for /dev/vgdemo has been saved in /etc/lvmconf/vgdemo.conf
```

Regardless of the method used above, issue the following command to validate the creation of the new Volume Group.

```
Type: vgdisplay –v /dev/vgdemo

--- Volume groups ---
VG Name                     /dev/vgdemo
VG Write Access             read/write
VG Status                   available
Max LV                      255
Cur LV                      0
Open LV                     0
Max PV                      16
Cur PV                      1
Act PV                      1
Max PE per PV               1919
```

```
VGDA                    2
PE Size (Mbytes)        16
Total PE                1919
Alloc PE                0
Free PE                 1919
Total PVG               0
Total Spare PVs         0
Total Spare PVs in use  0

   --- Physical volumes ---
   PV Name                     /dev/dsk/c27t0d0
   PV Status                   available
   Total PE                    1919
   Free PE                     1919
   Autoswitch                  On
```

Note

Prior to HP-UX 11i v3, there was no native multi-path I/O solution. On HP-UX 11i v2 or v1 hosts, HP instead provided a native failover solution named "**PVLinks**". In this scenario, each and every LUN presented to a HP-UX 11i v2 or v1 host would be visible via two (2) independent and unique device paths. Identifying the secondary device path would involve running the following commands and observing the output.

Type: `ioscan -fnC disk`

```
Class     I  H/W Path                    Driver   S/W State   H/W Type     Description
===================================================================
disk      4  0/0/2/0.6.0                 sdisk    CLAIMED     DEVICE       SEAGATE ST318404LC
                                /dev/dsk/c1t6d0   /dev/rdsk/c1t6d0
disk     24  0/12/0/0.206.2.255.0.0.0    sdisk    CLAIMED     DEVICE       COMPELNTCompellent Vol
                                /dev/dsk/c31t0d0   /dev/rdsk/c31t0d0
disk     25  0/12/0/0.206.3.255.0.0.0    sdisk    CLAIMED     DEVICE       COMPELNTCompellent Vol
                                /dev/dsk/c32t0d0   /dev/rdsk/c32t0d0
disk     26  0/12/0/0.206.4.255.0.0.0    sdisk    CLAIMED     DEVICE       COMPELNTCompellent Vol
                                /dev/dsk/c33t0d0   /dev/rdsk/c33t0d0
disk     27  0/12/0/0.206.5.255.0.0.0    sdisk    CLAIMED     DEVICE       COMPELNTCompellent
Vol
                                /dev/dsk/c34t0d0   /dev/rdsk/c34t0d0
disk      6  1/4/0/0.206.2.255.0.0.0     sdisk    CLAIMED     DEVICE       COMPELNTCompellent
Vol
                                /dev/dsk/c27t0d0   /dev/rdsk/c27t0d0
disk     28  1/4/0/0.206.3.255.0.0.0     sdisk    CLAIMED     DEVICE       COMPELNTCompellent
Vol
                                /dev/dsk/c28t0d0   /dev/rdsk/c28t0d0
disk     29  1/4/0/0.206.4.255.0.0.0     sdisk    CLAIMED     DEVICE       COMPELNTCompellent
Vol
```

```
                                        /dev/dsk/c29t0d0        /dev/rdsk/c29t0d0
disk    30  1/4/0/0.206.5.255.0.0.0     sdisk           CLAIMED     DEVICE      COMPELNTCompellent
Vol
                                        /dev/dsk/c30t0d0        /dev/rdsk/c30t0d0
```

As shown in RED above and for sake of this discussion, please assume that the device located at this address 1/4/0/0.206.2.255.0.0.0 is the primary path to the device /dev/dsk/c27t0d0. In order to determine the alternate path to this device, issue the command.

Type: `vgscan –pv`

(It is **very** important to note the use of **-p** parameter i.e. preview; the lack of the preview parameter would otherwise rebuild the /etc/lvmtab file with potentially harmful results).

```
vgscan: Warning: couldn't query physical volume "/dev/dsk/c27t0d0":
The specified path does not correspond to physical volume attached to
this volume group
vgscan: Warning: couldn't query all of the physical volumes.
vgscan: Warning: couldn't query physical volume "/dev/dsk/c27t0d1":
The specified path does not correspond to physical volume attached to
this volume group
vgscan: Warning: couldn't query physical volume "/dev/dsk/c27t0d2":
The specified path does not correspond to physical volume attached to
this volume group
<snip>
Couldn't stat physical volume "/dev/dsk/c29t0d0":
Invalid argument
Couldn't stat physical volume "/dev/dsk/c30t0d0":
Invalid argument

/dev/vg00
/dev/dsk/c1t6d0

/dev/vgdemo
/dev/dsk/c31t0d0
/dev/dsk/c27t0d0
```

(The warnings that show above are normal for the “vgscan” command; note the area shown in RED above, representing the two (2) paths identified, primary “c27t0d0” and alternate “c31t0d0”)

We now add the alternate path to the Volume Group with the command.

Type: `vgextend /dev/vgdemo /dev/dsk/c31t0d0`

```
Volume group "/dev/vgdemo" has been successfully extended.
Volume Group configuration for /dev/vgdemo has been saved in /etc/lvmconf/vgdemo.conf
```

To confirm that the alternate additional path to the LUN has been added to the Volume Group, issue the command.

Type: `vgdisplay –v /dev/vgdemo`

```
--- Volume groups ---
VG Name                     /dev/vgtest
VG Write Access             read/write
VG Status                   available
Max LV                      255
Cur LV                      1
Open LV                     1
Max PV                      16
Cur PV                      1
Act PV                      1
Max PE per PV               1919
VGDA                        2
PE Size (Mbytes)            16
Total PE                    1919
Alloc PE                    1919
Free PE                     0
Total PVG                   0
Total Spare PVs             0
Total Spare PVs in use      0

   --- Physical volumes ---
   PV Name                  /dev/dsk/c27t0d0
   PV Name                  /dev/dsk/c31t0d0 Alternate Link
   PV Status                available
   Total PE                 1919
   Free PE                  0
   Autoswitch               On
   Proactive Polling        On
```

Creating a new Logical Volume

At this point, we can now create a new Logical Volume. We will create a Logical Volume, which is 10GB in size. To achieve this, we issue the command.

Type: `lvcreate –L 10000 /dev/vgdemo`

(where **-L**: use 10000 in MB as the size of the new Logical Volume; using the –L parameter allows LVM to determine the number of Physical Extents to use in creating this Logical Volume)

```
Logical volume "/dev/vgdemo/lvol1" has been successfully created with
character device "/dev/vgdemo/rlvol1".
```

```
Logical volume "/dev/vgdemo/lvol1" has been successfully extended.
Volume Group configuration for /dev/vgdemo has been saved in /etc/lvmconf/vgdemo.conf
```

To validate the Logical Volume creation, issue the following command.

Type: `lvdisplay –v /dev/vgdemo/lvol1`

```
--- Logical volumes ---
LV Name                     /dev/vgdemo/lvol1
VG Name                     /dev/vgdemo
LV Permission               read/write
LV Status                   available/syncd
Mirror copies               0
Consistency Recovery        MWC
Schedule                    parallel
LV Size (Mbytes)            10240
Current LE                  640
Allocated PE                640
Stripes                     0
Stripe Size (Kbytes)        0
Bad block                   on
Allocation                  strict
IO Timeout (Seconds)        default
```

We can now prepare the Logical Volume for use by creating a filesystem on it. Issue the command.

Type: `newfs –F vxfs –o largefiles /dev/vgdemo/rlvol1`

(where **-o largefiles**: allows for any single files in the Logical Volume to be 2GB or larger)

```
version 6 layout
10256384 sectors, 10256384 blocks of size 1024, log size 16384 blocks
largefiles supported
```

Mounting and using the new Logical Volume/Filesystem

We will now mount this newly created Logical Volume to the /vgdemo directory. We create the new directory with the command.

Type: `mkdir –p /vgdemo`

We then add the following entry to the /etc/fstab file with the following command.

Type: `echo “/dev/vgdemo/lvol1 /vgdemo vxfs largefiles,delaylog 0 2” >> /etc/fstab`

(where **largefiles**: informs LVM to mount this filesystem with the largefiles option during and with every subsequent system reboot)

Now issue the following command to mount this new filesystem.

Type: `mount /vgdemo`

Confirm that the newly created and mounted filesystem is ready for use with the command.

Type: `bdf`

```
Filesystem          kbytes    used   avail %used Mounted on
/dev/vg00/lvol3       1048576  304200  738600   29% /
/dev/vg00/lvol1       1835008  295368 1527648   16% /stand
/dev/vg00/lvol8       4718592 1003216 3687448   21% /var
/dev/vg00/lvol7       3940352 3216616  718088   82% /usr
/dev/vg00/lvol4        524288  114136  406952   22% /tmp
/dev/vg00/lvol6       4964352 2888208 2059976   58% /opt
/dev/vg00/lvol5        106496   68656   37552   65% /home
/dev/vgdemo/lvol1    10256386   19595 9596997    0% /vgdemo
```

Dynamic LUN Expansion (DLE) and Dynamic LUN Contraction (DLC)

This section builds upon the example used in the previous sections. Please also familiarize yourself with the HP documentation about the new "'vgmodify" command before attempting this procedure on any production system.

- Using the vgmodify command to perform LVM Volume Group DLE and DLC
  http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01920387/c01920387.pdf

With HP-UX 11i v2, the vgmodify command is available with the application of patch PHCO_35524 or later. Please reference your respective support channels via the HP Support Center to determine the latest stable patch which may have superseded PHCO_35524. In addition, the vgmodify command on HP-UX 11i v2 **DOES NOT** support the –C parameter (i.e. it lacks the ability to perform DLC).

Expanding the LUN

As you see below, we once again that the 30GB LUN as presented in previous examples.

```
Class     I  H/W Path                    Driver   S/W State   H/W Type     Description
===================================================================================
disk      4  0/0/2/0.6.0                 sdisk    CLAIMED     DEVICE       SEAGATE ST318404LC
                                         /dev/dsk/c1t6d0       /dev/rdsk/c1t6d0
<snip>
disk      6  1/4/0/0.206.2.255.0.0.0     sdisk    CLAIMED     DEVICE       COMPELNTCompellent Vol
                                         /dev/dsk/c27t0d0      /dev/rdsk/c27t0d0
disk     28  1/4/0/0.206.3.255.0.0.0     sdisk    CLAIMED     DEVICE       COMPELNTCompellent Vol
                                         /dev/dsk/c28t0d0      /dev/rdsk/c28t0d0
disk     29  1/4/0/0.206.4.255.0.0.0     sdisk    CLAIMED     DEVICE       COMPELNTCompellent Vol
                                         /dev/dsk/c29t0d0      /dev/rdsk/c29t0d0
disk     30  1/4/0/0.206.5.255.0.0.0     sdisk    CLAIMED     DEVICE       COMPELNTCompellent Vol
                                         /dev/dsk/c30t0d0      /dev/rdsk/c30t0d0
```

On the OS, the geometry of this LUN is shown below with the command.

```
Type: /etc/vx/diag.d/vxscsi –g /dev/rdsk/c27t0d0

geometry of /dev/rdsk/c27t0d0: nhead=8 nsect=158 ncyl=49774 sectsz=512 rpm=15000 cap=62914560
```

(where 62914560 x 512KB ~ 30GB LUN)

We are now going to attempt to expand this 30GB LUN to 50GB in size. On the Dell Compellent Storage Center, locate the LUN in the "Storage -> Volumes" tree. We will the right-Click on the LUN and choose the "Expand Volume" option. We will then define the new size for the LUN (in this case 50GB) then left-Click on the "Expand Volume Now" button.

Running the "vxscsi" command once more, you will notice that the OS has detected the geometry change of the LUN immediately, as shown in RED below.

```
Type: /etc/vx/diag.d/vxscsi –g /dev/rdsk/c27t0d0
```

geometry of /dev/rdsk/c27t0d0: nhead=8 nsect=263 ncyl=49837 sectsz=512 rpm=15000 cap=104857600
(where 104857600 x 512KB ~ 50GB LUN)

If this LUN is NOT presently part of a Volume Group, you can now create a Volume Group with the procedures located here Creating a new Volume Group.

If this LUN is already part of a Volume Group, you would now want to expose the additional capacity of the LUN to the Volume Group that it is a part of. This is where the "vgmodify" commands come into use.

Exposing the new capacity to the Volume Group

Prior to expanding the LUN in the Dell Compellent Storage Center, your Volume Group "vgdemo" would have looked like this; highlighted in RED to reflect current number of PEs (Physical Extents) of the LUN.

```
--- Volume groups ---
VG Name                     /dev/vgdemo
VG Write Access             read/write
VG Status                   available
Max LV                      255
Cur LV                      0
Open LV                     0
Max PV                      16
Cur PV                      1
Act PV                      1
Max PE per PV               1919
VGDA                        2
PE Size (Mbytes)            16
Total PE                    1919
Alloc PE                    0
Free PE                     1919
Total PVG                   0
Total Spare PVs             0
Total Spare PVs in use      0


   --- Physical volumes ---
   PV Name                     /dev/dsk/c27t0d0
   PV Status                   available
   Total PE                    1919
   Free PE                     1919
   Autoswitch                  On
   Proactive Polling           On
```

Though the OS had detected the geometry changes of the LUN (see Expanding the LUN), the Volume Group to which it belongs is not yet aware of this change. With HP-UX 11i v2, the process to increase the number of available PEs (Physical Extents) of the LUN within the Volume Group (and thus also increase the total available capacity of the Volume Group), would need to be performed in **OFFLINE**

mode. In any production environment, this would also require the necessary planning for downtime, as well as getting any necessary change management approvals.

The first step in this process is to ensure that the Volume Group is OFFLINE (inactive) and if it is part of a cluster (MC/Serviceguard or otherwise), then also made cluster unaware; if PE/Renumbering is required, it would also need to be addressed accordingly. For the sake of simplicity in this example, we shall assume that this Volume Group is NOT part of a cluster nor is PE/Renumbering required. For further details about either of these latter two topics, refer to this document below.

- Using the vgmodify command to perform LVM Volume Group DLE and DLC
  http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01920387/c01920387.pdf

Now we need to determine the appropriate configuration parameters for use with the "vgmodify" command. While the Volume Group is still in an ONLINE (active) state, issue the command.

Type: `vgmodify -t -v vgdemo`

```
Volume Group configuration for /dev/vgdemo has been saved in /etc/lvmconf/vgdemo.conf

Current Volume Group settings:
                                Max LV      255
                                Max PV      16
                            Max PE per PV   1919
                           PE Size (Mbytes)    16
                           VGRA Size (Kbytes)   304

VGRA space (Kbytes) on all Physical Volumes:
          PV                        current        -n
      /dev/rdsk/c27t0d1                1088      16384
      Summary                          1088      16384

Volume Group optimized settings (no PEs renumbered):
  max_pv(-p)      max_pe(-e)      Disk size (Mb)
  1               65535           1048561
  2               65532           1048513
  3               43772           700353
  4               32764           524225
  5               26108           417729
  6               21756           348097
  7               18684           298945
  8               16380           262081
  9               14588           233409
  10              13052           208833
  11              11772           188353
  12              10748           171969
<snip>
```

From this output above, we would want to identify the most optimized set of parameters based upon approximated LUN sizes, which would be added to this Volume Group. Since the initial LUN was 30GB in size (and since expanded to 50GB on the Dell Compellent Storage Center), and with the assumption that subsequent LUNs will max out around 250GB, we will use the parameter set identified in RED above. This parameter set will reconfigure the Volume Group to accommodate a maximum of 8 PVs (Physical Volumes i.e. LUNs) with each LUN having 16380 PEs (Physical Extents) of 16MB (PE Size) each. We would preview this parameter set with the command below.

Type: `vgmodify -r -p 8 -e 16380 vgdemo`

```
Current Volume Group settings:
                    Max LV      255
                    Max PV       16
                Max PE per PV    1919
              PE Size (Mbytes)     16
            VGRA Size (Kbytes)    304
The current and new Volume Group parameters differ.
An update to the Volume Group IS required


New Volume Group settings:
                    Max LV      255
                    Max PV        8
                Max PE per PV   16380
              PE Size (Mbytes)     16
            VGRA Size (Kbytes)   1088
Review complete. Volume group not modified
```

To effect this change, we would first need to OFFLINE (inactive) the Volume Group, issue the command.

Type: `vgchange –a n vgdemo`

```
Volume group "vgdemo" has been successfully changed.
```

We would then issue the “vgmodify” command from above. We would run it this time without the “**-r**” (ie. preview) parameter.

Type: `vgmodify -p 8 -e 16380 vgdemo`

```
Current Volume Group settings:
                    Max LV      255
                    Max PV       16
                Max PE per PV    1919
              PE Size (Mbytes)     16
            VGRA Size (Kbytes)    304
The current and new Volume Group parameters differ.
```

```
"/dev/rdsk/c27t0d1" size changed from 31457280 to 52428800kb
An update to the Volume Group IS required


New Volume Group settings:
                                Max LV      255
                                Max PV      8
                            Max PE per PV   16380
                          PE Size (Mbytes)      16
                         VGRA Size (Kbytes)   1088
New Volume Group configuration for "vgdemo" has been saved in "/etc/lvmconf/vgdemo.conf"
Old Volume Group configuration for "vgdemo" has been saved in "/etc/lvmconf/vgdemo.conf.old"
Starting the modification by writing to all Physical Volumes
Applying the configuration to all Physical Volumes from "/etc/lvmconf/vgdemo.conf"
```

Finally, we would bring the Volume Group ONLINE (active) again.

Type: `vgchange -a y vgdemo`

```
Volume group "vgdemo" has been successfully changed.
```

Checking the Volume Group once more, we would now see the output below. The process of expanding any Logical Volumes (and its accompanying filesystem via lvextend, fsadm, vxresize etc. and any respective and/or necessary licensing of such tools) which may reside on this Volume Group is outside the scope of this document; kindly refer to the HP Support Center for further instructions and/or assistance.

```
<snip>


   --- Physical volumes ---
   PV Name                  /dev/dsk/c27t0d0
   PV Status                available
   Total PE                 3199
   Free PE                  3199
   Autoswitch               On
   Proactive Polling        On
```

HP provides no tools/support nor facilities to accommodate online DLE/DLC on HP-UX 11i v1 or older platforms. Any LUN or Volume Group expansion efforts will need to be performed in a completely OFFLINE mode. The procedures to achieve this are outside the scope of this document; refer to the HP Support Center for further instructions and/or assistance.

Dynamic LUN Contraction (DLC) and Thoughts

With HP-UX 11i v2 (and the patch PHCO_35524 or later), the "vgmodify" command DOES NOT provide a "-C" parameter for the resizing of a PV (Physical Volume). Any LUN (PV), Logical Volume or Volume Group reduction effort will need to be performed in a completely OFFLINE mode. The procedures to achieve this are outside the scope of this document; kindly refer to the HP Support Center for further instructions and/or assistance.

It is **strongly recommended** as a manner of best practice to capture and maintain a known good backup copy of your data **PRIOR** to performing any LV/PV reduction procedures. The procedures to achieve this are outside the scope of this document; refer to the HP Support Center for further instructions and/or assistance.

It is also to be noted, that the Dell Compellent Storage Center does not presently support (nor provide the facility to perform) the reduction of LUN sizing on the Dell Compellent Storage Center. It is recommend that any data migration or relocation efforts be performed by presenting a secondary (and smaller) LUN to the target host, prepping the LUN for LVM use and relocating the data via OS-based tools (pvmove, tar, dd etc.); thereafter retiring and reclaiming the larger and unused LUN. The procedures to achieve this are outside the scope of this document; refer to the HP Support Center for further instructions and/or assistance.

With HP-UX 11i v1 and older, HP does not provide tools or support nor facilities to accommodate online DLE/DLC. Any LUN (PV), Logical Volume or Volume Group expansion/reduction effort will need to be performed in a completely OFFLINE mode. The procedures to achieve this are outside the scope of this document; kindly refer to the HP Support Center for further instructions and/or assistance.

Data Instant Replay (Continuous Snapshots)

It is often required in an enterprise environment to freeze and capture the state of a data set at a specific point in time; there are many and various reasons that this is needed, including but not limited to data backups, configuration snapshots prior to system upgrades etc. On the Dell Compellent Storage Center products, this facility and ability to capture snapshots is known as "Data Instant Replay" or "replay/s" for short. We will be using the term "replay" for brevity in this section and further discussion.

Prior to capturing a replay, it is recommended that the data set be put into a known and consistent state. This often and typically means that any writes to the data set (both buffered and direct) are flushed and the data set is frozen for the duration when the replay is captured. This is often known in the technical realm as "quiesceing" the data. For the sake of simplicity, we will be recreating a Volume Group named "vgdemo" with a single 100g LUN from the Dell Compellent Storage Center and a single 50g Logical Volume therefrom. The Volume Group and mounted filesystem looks as follow.

Type: `vgdisplay –v vgdemo`

```
--- Volume groups ---
VG Name                     /dev/vgdemo
VG Write Access             read/write
VG Status                   available
Max LV                      255
Cur LV                      1
Open LV                     1
Max PV                      16
Cur PV                      1
Act PV                      1
Max PE per PV               6399
VGDA                        2
PE Size (Mbytes)            16
Total PE                    6399
Alloc PE                    3125
Free PE                     3274
Total PVG                   0
Total Spare PVs             0
Total Spare PVs in use      0

   --- Logical volumes ---
   LV Name                  /dev/vgdemo/lvol1
   LV Status                available/syncd
   LV Size (Mbytes)         50000
   Current LE               3125
   Allocated PE             3125
   Used PV                  1


   --- Physical volumes ---
```

```
   PV Name                /dev/dsk/c11t0d1
   PV Name                /dev/dsk/c47t0d1 Alternate Link
   PV Status              available
   Total PE               6399
   Free PE                3274
   Autoswitch             On
   Proactive Polling      On
```

Type: `bdf`

```
Filesystem          kbytes    used   avail %used Mounted on
/dev/vg00/lvol3     1048576  322096  720864   31% /
/dev/vg00/lvol1     2097152   85800 1995744    4% /stand
/dev/vg00/lvol8     6029312 2174176 3826928   36% /var
/dev/vg00/lvol7     6914048 2663360 4217488   39% /usr
/dev/vg00/lvol6     4194304   16824 4144848    0% /tmp
/dev/vg00/lvol5     10518528 4065208 6402936   39% /opt
/dev/vg00/lvol4      524288   16696  503632    3% /home
rx1620.techsol.beer.town:/Tools
                    51380224 15007672 34099344   31% /Tools
/dev/vgdemo/lvol1   51200000   29634 47972226    0% /vgdemo/lvol1
```

In larger and more complex Volume Groups (i.e. for the purposes of database use or otherwise), it may become necessary to capture the replay of multiple LUNs on the Dell Compellent Storage Center at the same exact point in time. This is accomplished using a facility known as Consistency Groups. The discussion of Consistency Groups is outside the scope of this document and will be covered in a separate paper.

Quiesce the Volume Group

With HP-UX 11i v2 or older, you would need to unmount all potentially impacted filesystems and then disable the Volume Group.

Type: `umount /vgdemo/lvol1; bdf`

```
Filesystem          kbytes    used   avail %used Mounted on
/dev/vg00/lvol3     1048576  322096  720864   31% /
/dev/vg00/lvol1     2097152   85800 1995744    4% /stand
/dev/vg00/lvol8     6029312 2174176 3826928   36% /var
/dev/vg00/lvol7     6914048 2663360 4217488   39% /usr
/dev/vg00/lvol6     4194304   16824 4144848    0% /tmp
/dev/vg00/lvol5     10518528 4065208 6402936   39% /opt
/dev/vg00/lvol4      524288   16696  503632    3% /home
rx1620.techsol.beer.town:/Tools
                    51380224 15007672 34099344   31% /Tools
```

```
Type: vgchange –a n /dev/vgdemo
```

```
Volume group "/dev/vgdemo" has been successfully changed.
```

You would now capture the replay of this LUN on the Dell Compellent Storage Center (as will be discussed in the section below). When complete you would reenable I/O (writes) to the Volume Group with the command “vgchange –a y /dev/vgdemo”.

Capturing the Snapshot

With the Volume Group quisced, we would now identify the LUN on the Dell Compellent Storage Center and capture a replay of this LUN.

In this case, the LUN has been identified as “DanPreTest106” as matched by its serial number. We would right-Click on the LUN, select the “Replay” link and then the “Create Replay” option.

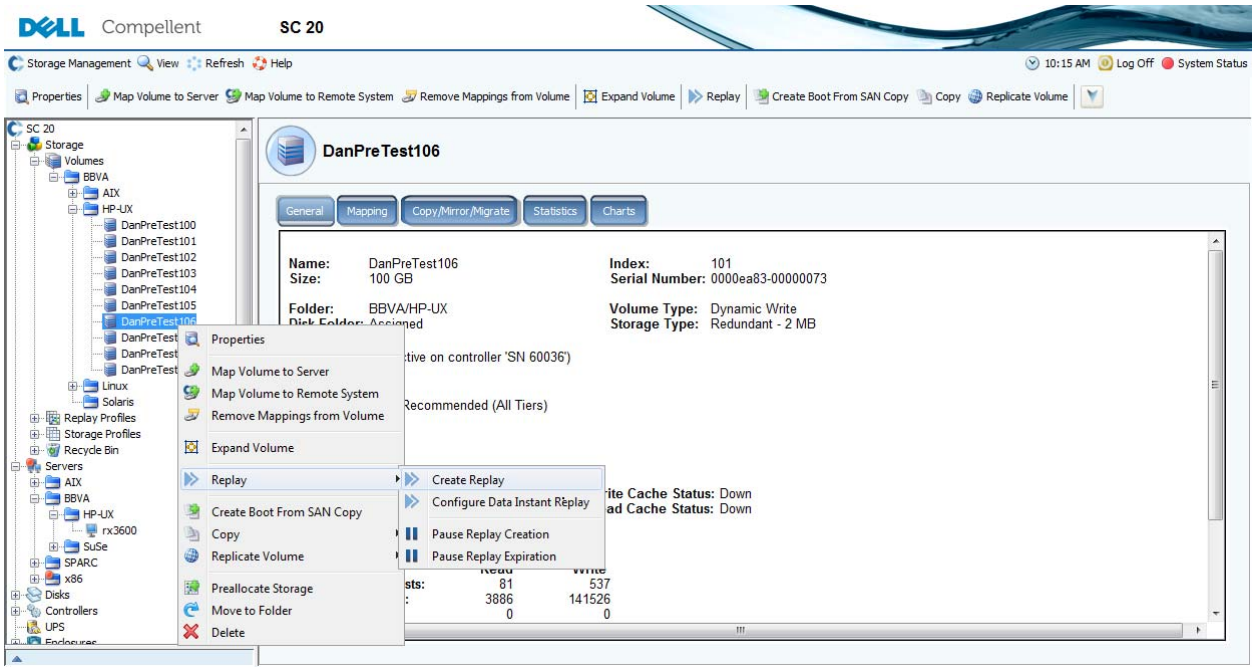

In this subsequent dialog, we have chosen to expire this replay in 4hrs (user configurable) and named the replay “Replay Jan 2012” (user definable field). We now left-click on the “Create Now” button to capture the replay.

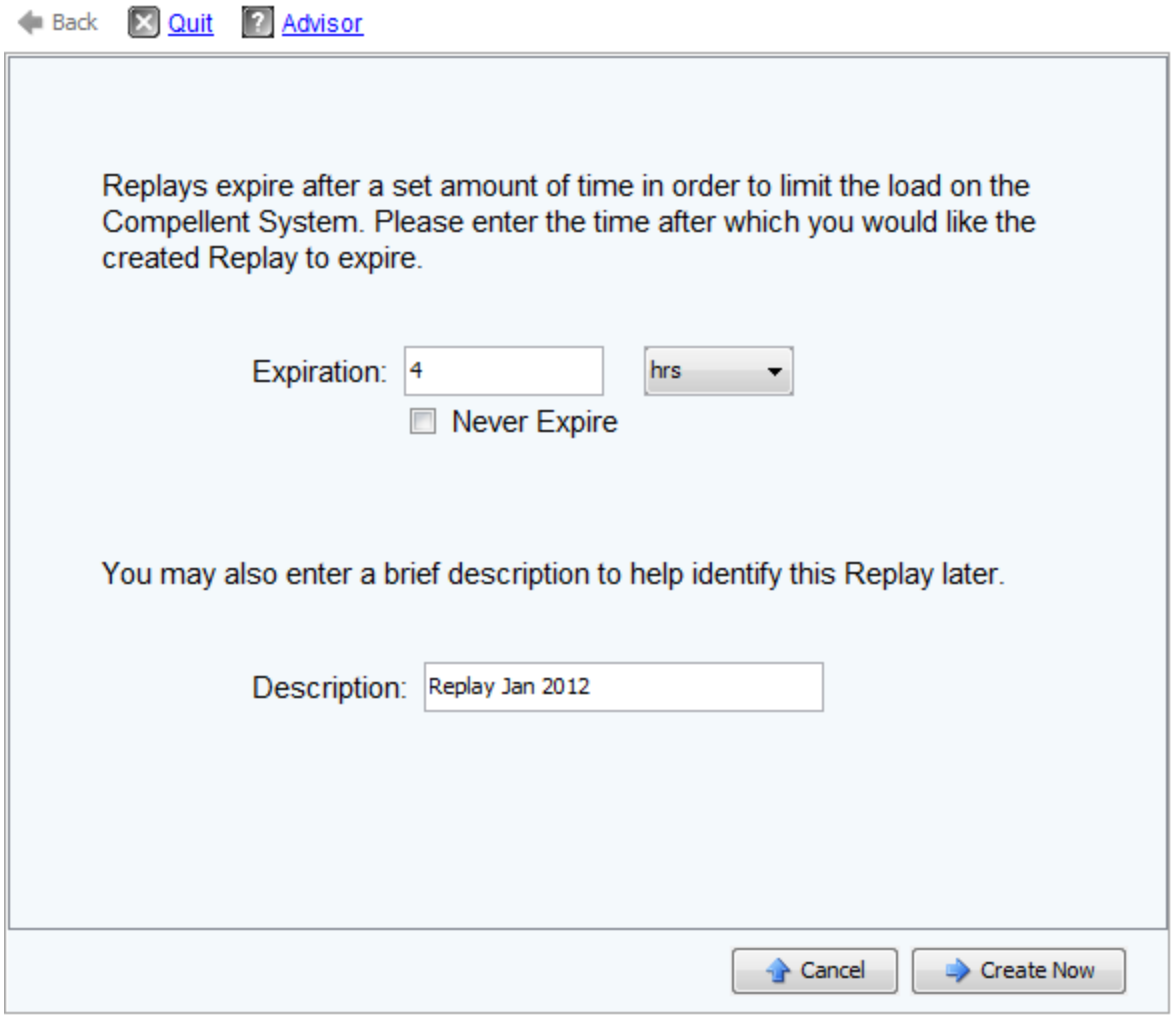
Back
Quit
Advisor
Replays expire after a set amount of time in order to limit the load on the Compellent System. Please enter the time after which you would like the created Replay to expire.
Expiration:
4
hrs
Never Expire
You may also enter a brief description to help identify this Replay later.
Description:
Replay Jan 2012
Cancel
Create Now

Back in the main Dell Compellent Storage Center window, you will notice a new button named "Replays" across the top button bar, left-Click to select it.

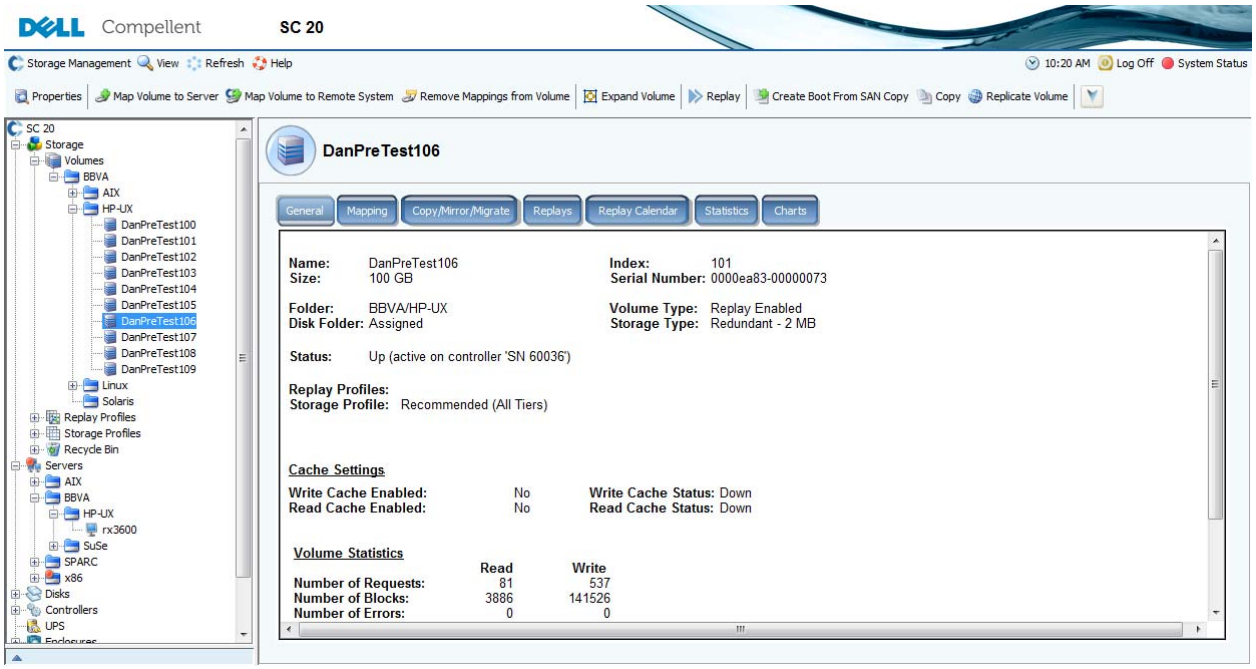

You will notice that the replay just captured in highlighted in blue.

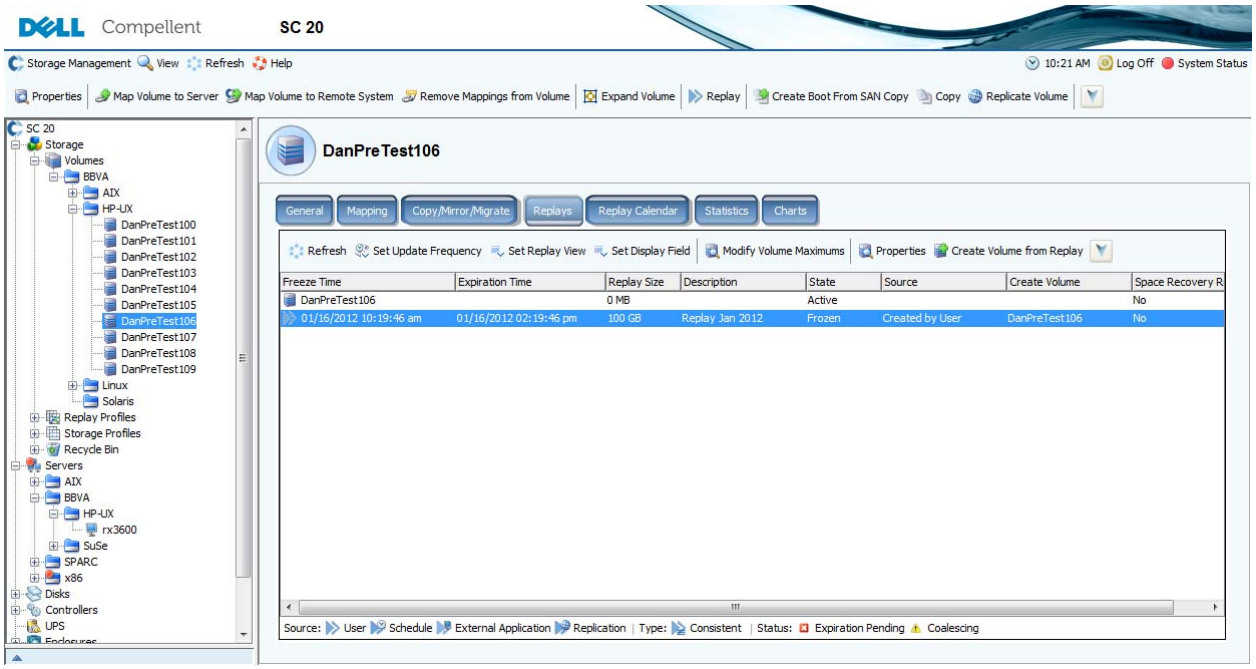

The Dell Compellent Storage Center replays cannot directly be remounted back to its original or alternate host. You would need to create a new Volume/LUN from this replay and then map this new Volume/LUN to the intended host; right-Click on the replay and select the “Create Volume from Replay” option.

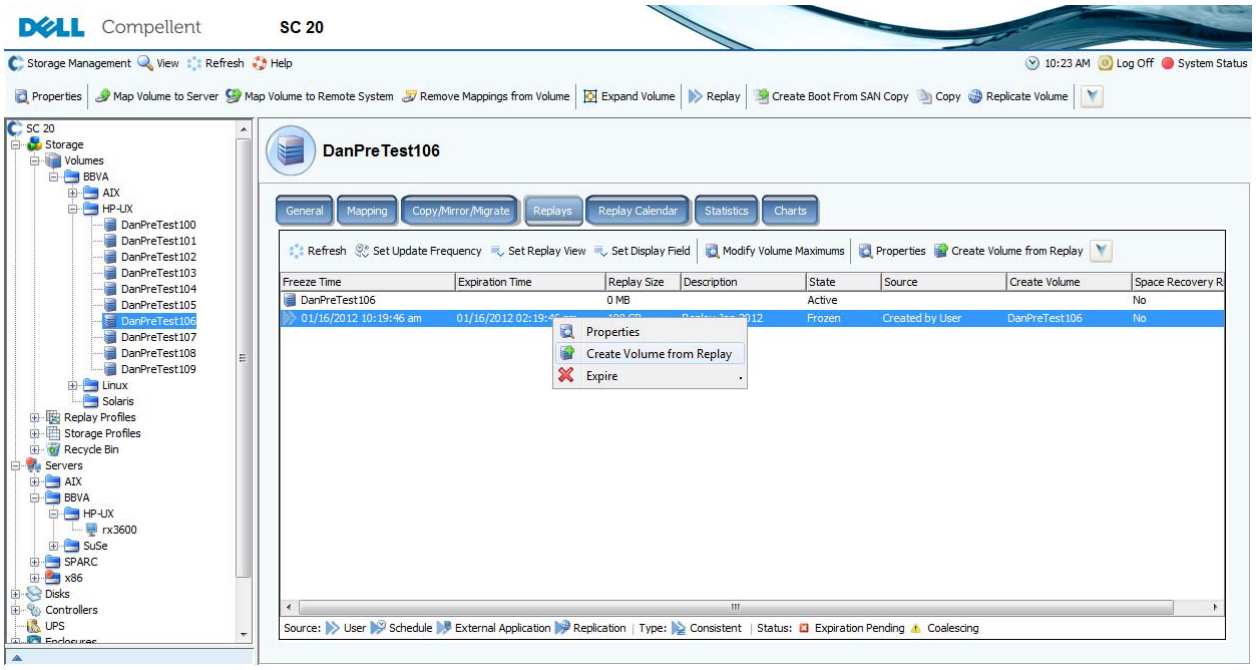

In the next dialog, give the new Volume/LUN a new name. In this case we will accept the default name being "DanPreTest106 View 1", then left-Click the "Create Now" button, then left-Click the "Quit" option when presented with the "Map Volume to Server" dialog. As you can see, the new Volume "DanPreTest 106 View 1" is now shown in the screenshot below in an unmapped state.

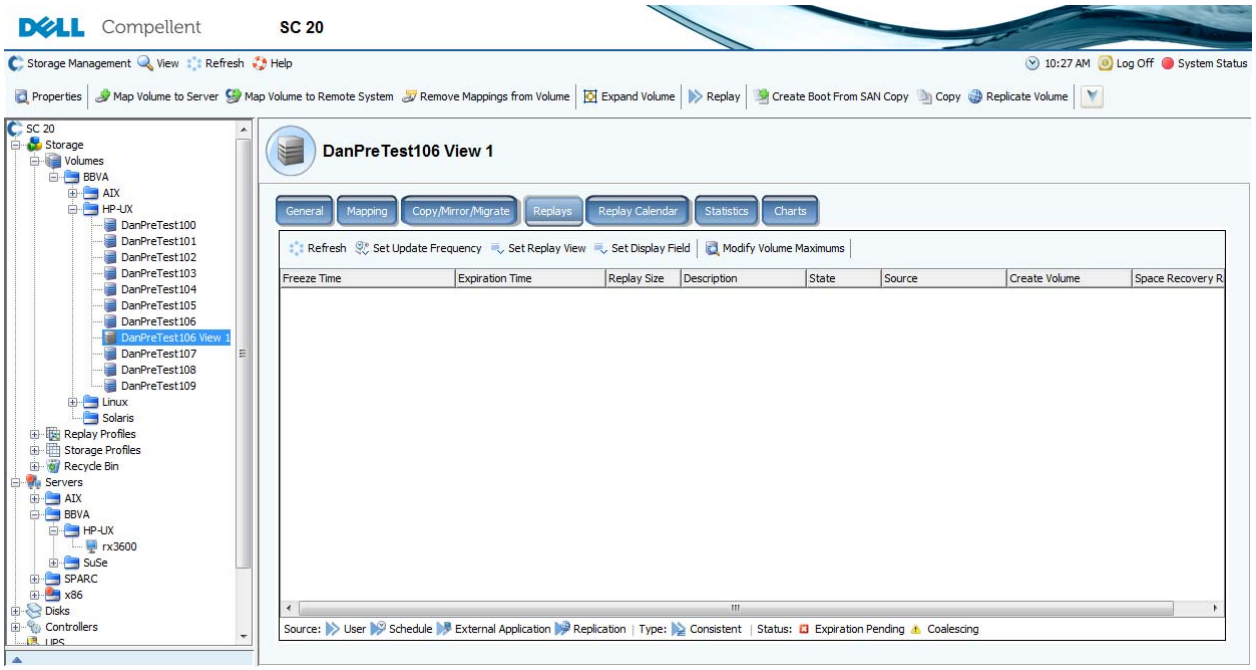

Mounting to the same host

Prior to mapping this LUN to a host, capture the following output from the host STDOUT.

Type: `ioscan -kfnC disk`

```
<snip>
disk       7  0/2/1/0/4/0.236.1.0.0.0.1  sdisk        CLAIMED    DEVICE      COMPELNTCompellent Vol
                      /dev/dsk/c11t0d1   /dev/rdsk/c11t0d1
disk      12  0/2/1/0/4/1.236.9.0.0.0.1  sdisk        CLAIMED    DEVICE      COMPELNTCompellent Vol
                      /dev/dsk/c47t0d1   /dev/rdsk/c47t0d1
```

Type: `strings /etc/lvmtab`

```
/dev/vg00
/dev/dsk/c57t0d0s2
/dev/vgdemo
/dev/dsk/c11t0d1
/dev/dsk/c47t0d1
```

Type: `vgexport -p -v -m /tmp/vgdemo.map vgdemo`

```
Beginning the export process on Volume Group "vgdemo".
/dev/dsk/c11t0d1
/dev/dsk/c47t0d1
```

Please refer to the section Server Setup for instructions on mapping a LUN from the Dell Compellent Storage Center to a host. In this case, we will be mapping “DanPreTest106 View 1” back to the same host.

After LUN representation, run the following commands.

Type: `insf -e; ioscan -fnC disk`

```
<snip>
disk       7  0/2/1/0/4/0.236.1.0.0.0.1  sdisk        CLAIMED    DEVICE      COMPELNTCompellent Vol
                      /dev/dsk/c11t0d1   /dev/rdsk/c11t0d1
disk       8  0/2/1/0/4/0.236.1.0.0.0.2  sdisk        CLAIMED    DEVICE      COMPELNTCompellent Vol
                      /dev/dsk/c11t0d2   /dev/rdsk/c11t0d3
disk      12  0/2/1/0/4/1.236.9.0.0.0.1  sdisk        CLAIMED    DEVICE      COMPELNTCompellent Vol
                      /dev/dsk/c47t0d1   /dev/rdsk/c47t0d1
disk      13  0/2/1/0/4/1.236.9.0.0.0.2  sdisk        CLAIMED    DEVICE      COMPELNTCompellent Vol
                      /dev/dsk/c47t0d2   /dev/rdsk/c47t0d3
```

You will notice the newly presented LUN as shown in RED.

Since this new LUN shares the same Volume Group name and Volume Group Instance number as "vgdemo"; we will need to create a new Volume Group and change the Instance number on this LUN.

Type: `ls –la /dev/*/group | awk '{print $6}'`

(This tells us that only VGID 0 and 1 are in use, we will create the next Volume Group with VGID 2).

0x000000
0x010000

Type: `mkdir –p /dev/vgdemo_import; mknod /dev/vgdemo_import/group c 64 0x020000`

Type: `vgchgid /dev/rdsk/c11t0d3`

We would finally import the new LUN with the following command.

Type: `vgimport –v vgdemo_import /dev/dsk/c11t0d3`

Warning: A backup of this volume group may not exist on this machine.
Please remember to take a backup using the vgcfgbackup command after activating the volume group.

Activate this new Volume Group.

Type: `vgchange –a y /dev/vgdemo_import`

Activated volume group.
Volume group "/dev/vgdemo_import" has been successfully changed.

And extend the Volume Group with the alternate path to this LUN.

Type: `vgextend /dev/vgdemo_import /dev/dsk/c47t0d3`

Volume group "/dev/vgdemo_import" has been successfully extended.
Volume Group configuration for /dev/vgdemo_import has been saved in
/etc/lvmconf/vgdemo_import.conf

As a manner of best practice, we will then "fsck" this new Logical Volume and mount it to a new directory path. This would conclude the procedure in bringing this Volume back to the same host.

Type: `fsck –F vxfs /dev/vgdemo_import/lvol1`

file system is clean - log replay is not required

Type: `mkdir –p /vgdemo_import/lvol1`

Type: `mount –F vxfs –o largefiles /dev/vgdemo_import/lvol1 /vgdemo_import/lvol1`

Type: bdf

```
Filesystem          kbytes    used   avail %used Mounted on
/dev/vg00/lvol3    1048576  242712  799608   23% /
/dev/vg00/lvol1    2097152  175408 1906824    8% /stand
/dev/vg00/lvol8    4194304 1379368 2799440   33% /var
/dev/vg00/lvol7    6291456 3018472 3247424   48% /usr
/dev/vg00/lvol6    4194304   20928 4140776    1% /tmp
/dev/vg00/lvol5    10485760 5247280 5197656   50% /opt
/dev/vg00/lvol4     524288   20800  499560    4% /home
rx1620.techsol.beer.town:/Tools
                   51380224 11460912 37424433   23% /Tools
/dev/vgdemo/lvol1  51249152   79205 47971833    0% /vgdemo/lvol1
/dev/vgdemo_import/lvol1
                   51249152   79205 47971833    0% /vgdemo_import/lvol1
```

Mounting to a alternate host

Most of the procedures here are identical to the section Mounting to the same host above; with the exception of the following.

We would create a new Volume Group to land this newly imported LUN; give this new Volume Group a VGID, which does not conflict with any VGIDs already in use on this target host. We would also move the /tmp/vgdemo.map file from the originating host to /tmp on this target host. We would finally identify the device file of this new LUN on the target host. For the sake of discussion, we will assume that this new LUN has been discovered as /dev/dsk/c29t0d2 (/dev/rdsk/c29t0d2 accordingly). We will then import the new LUN with the command.

Type: vgimport –m /tmp/vgdemo.map –v vgdemo_import /dev/dsk/c29t0d2

```
Beginning the import process on Volume Group "vgdemo_import".
vgimport: Warning:  Volume Group belongs to different CPU ID.
Can not determine if Volume Group is in use on another system. Continuing.
Logical volume "/dev/vgdemo_import/lvol1" has been successfully created with lv number 1.
vgimport: Volume group "/dev/vgdemo_import" has been successfully created.
Warning: A backup of this volume group may not exist on this machine.
Please remember to take a backup using the vgcfgbackup command after activating the volume group.
```

We then vgchange to activate this new Volume Group vgdemo_import, fsck the new Logical Volume and mount the filesystem for use. The resulting output looks as follows.

Type: bdf

```
Filesystem          kbytes    used   avail %used Mounted on
/dev/vg00/lvol3    1048576  242712  799608   23% /
/dev/vg00/lvol1    2097152  175408 1906824    8% /stand
/dev/vg00/lvol8    4194304 1379368 2799440   33% /var
/dev/vg00/lvol7    6291456 3018472 3247424   48% /usr
```

```
/dev/vg00/lvol6    4194304   20928 4140776    1% /tmp
/dev/vg00/lvol5    10485760 5247280 5197656   50% /opt
/dev/vg00/lvol4     524288   20800  499560    4% /home
rx1620.techsol.beer.town:/Tools
                   51380224 11460912 37424433   23% /Tools
/dev/vgdemo/lvol1  51249152   79205 47971833    0% /vgdemo/lvol1
/dev/vgdemo_import/lvol1
                   51249152   79205 47971833    0% /vgdemo_import/lvol1
```

There are some very important caveats to note. HP-UX 11i v2 and older use LVM 1.0 structures while HP-UX 11i v3 uses LVM 1.0, 2.0 and/or 2.1 structures. **LVM structures are NOT backward compatible!** You could technically create a Volume Group and Logical Volume on HP-UX 11i v3 and access it with HP-UX 11i v1; however, **VxFS versions are also NOT backward compatible**. Therefore, even if you create an LVM 1.0 Volume Group with HP-UX 11i v3, the file system that you place on any Logical Volumes with HP-UX 11i v3 will not be readable by any previous version of HP-UX. **For all practical purposes, replay/s are OS forward compatible but not OS backward compatible! This is a limitation of the HP-UX, LVM structures and VxFS and not of Dell Compellent Storage Center replay functionality.**

Working with Striped/Mirror or MultiVolume Volume Groups

Most of the procedures here are identical to the section Mounting to the same host or Mounting to a alternate host above; with the exception of the following. For this example, we will be using a 5-way striped Volume Group; mirrored or concatenated Volume Groups would be managed in a similar manner.

To ensure the data integrity of these MultiVolume Volume Groups, we would use a feature of the Dell Compellent Storage Center known as a Consistency Group. A Consistency Group allows the user to capture a Replay of multiple volumes on the Dell Compellent Storage Array at the same and immediate point in time. Achieving this function would involve the use of Replay Profiles.

The first step is to create a Consistent Replay profile as seen below. The Replay Profiles sub-item is located within the Storage tree of the Dell Compellent Storage Center GUI. Right-Clicking on the Replay Profiles item displays this dialog shown below, where you would then select the “Create Consistent Replay Profile” option.

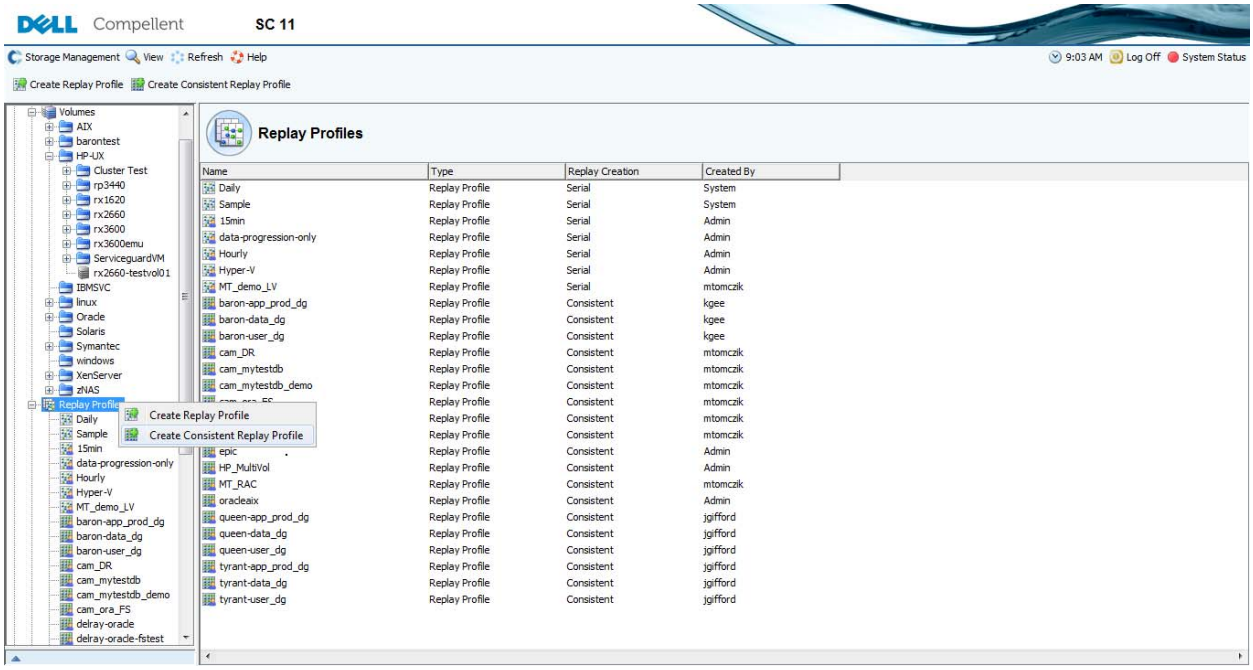

In the next dialog, read the Disclaimer then left-Click the “Continue” button.

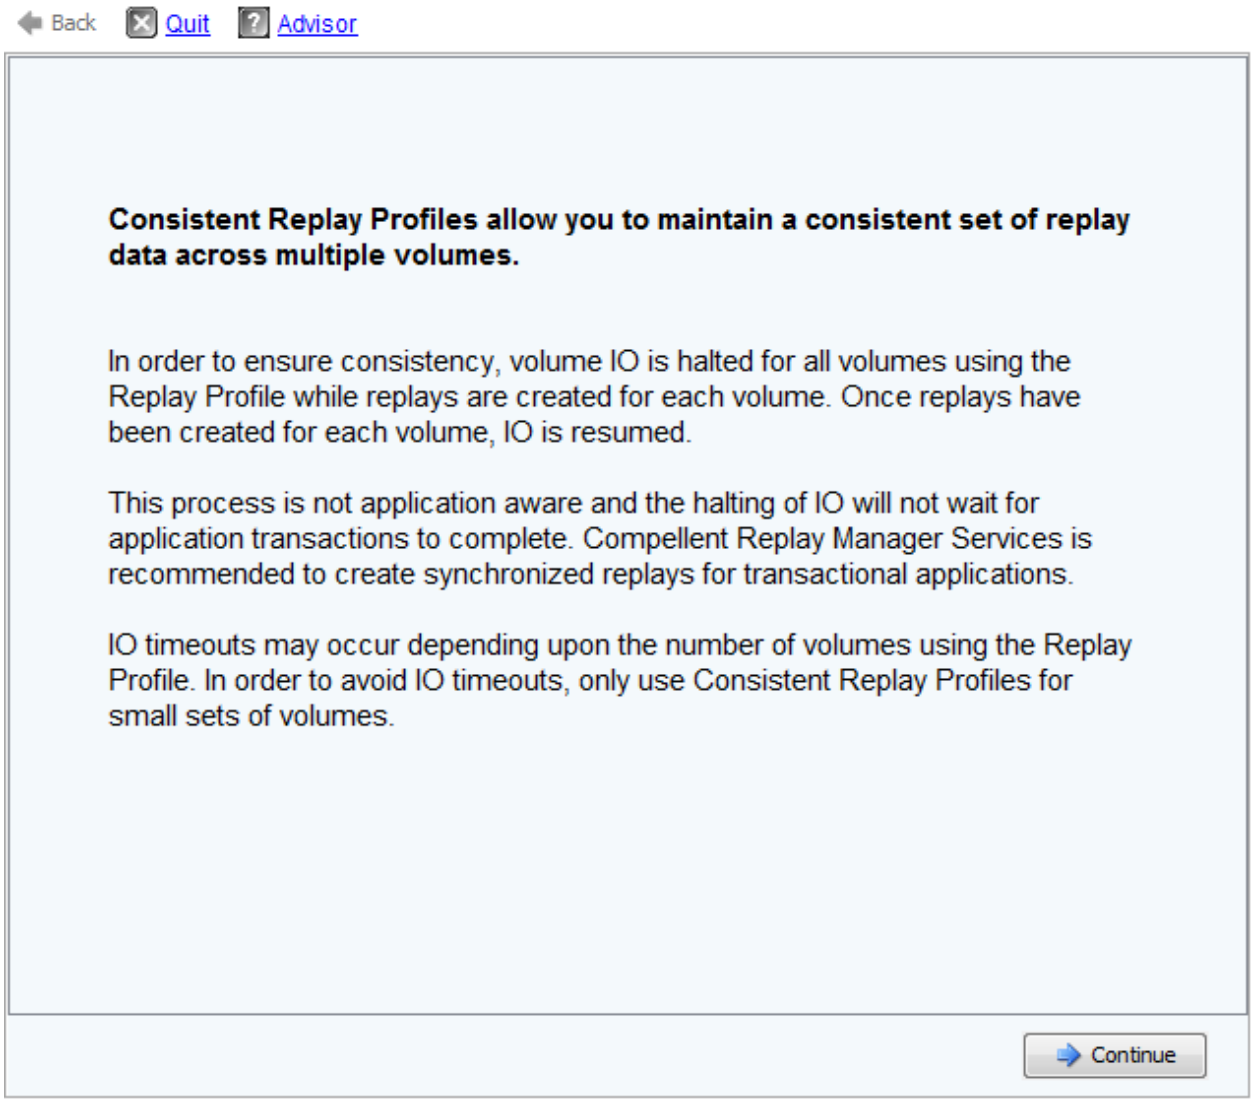
Back
Quit
Advisor
Consistent Replay Profiles allow you to maintain a consistent set of replay data across multiple volumes.
In order to ensure consistency, volume IO is halted for all volumes using the Replay Profile while replays are created for each volume. Once replays have been created for each volume, IO is resumed.
This process is not application aware and the halting of IO will not wait for application transactions to complete. Compellent Replay Manager Services is recommended to create synchronized replays for transactional applications.
IO timeouts may occur depending upon the number of volumes using the Replay Profile. In order to avoid IO timeouts, only use Consistent Replay Profiles for small sets of volumes.
Continue

In the next dialog, give your Consistent Replay profile a name (in this case, Demo_Profile) then left-Click the "Create Now" button.

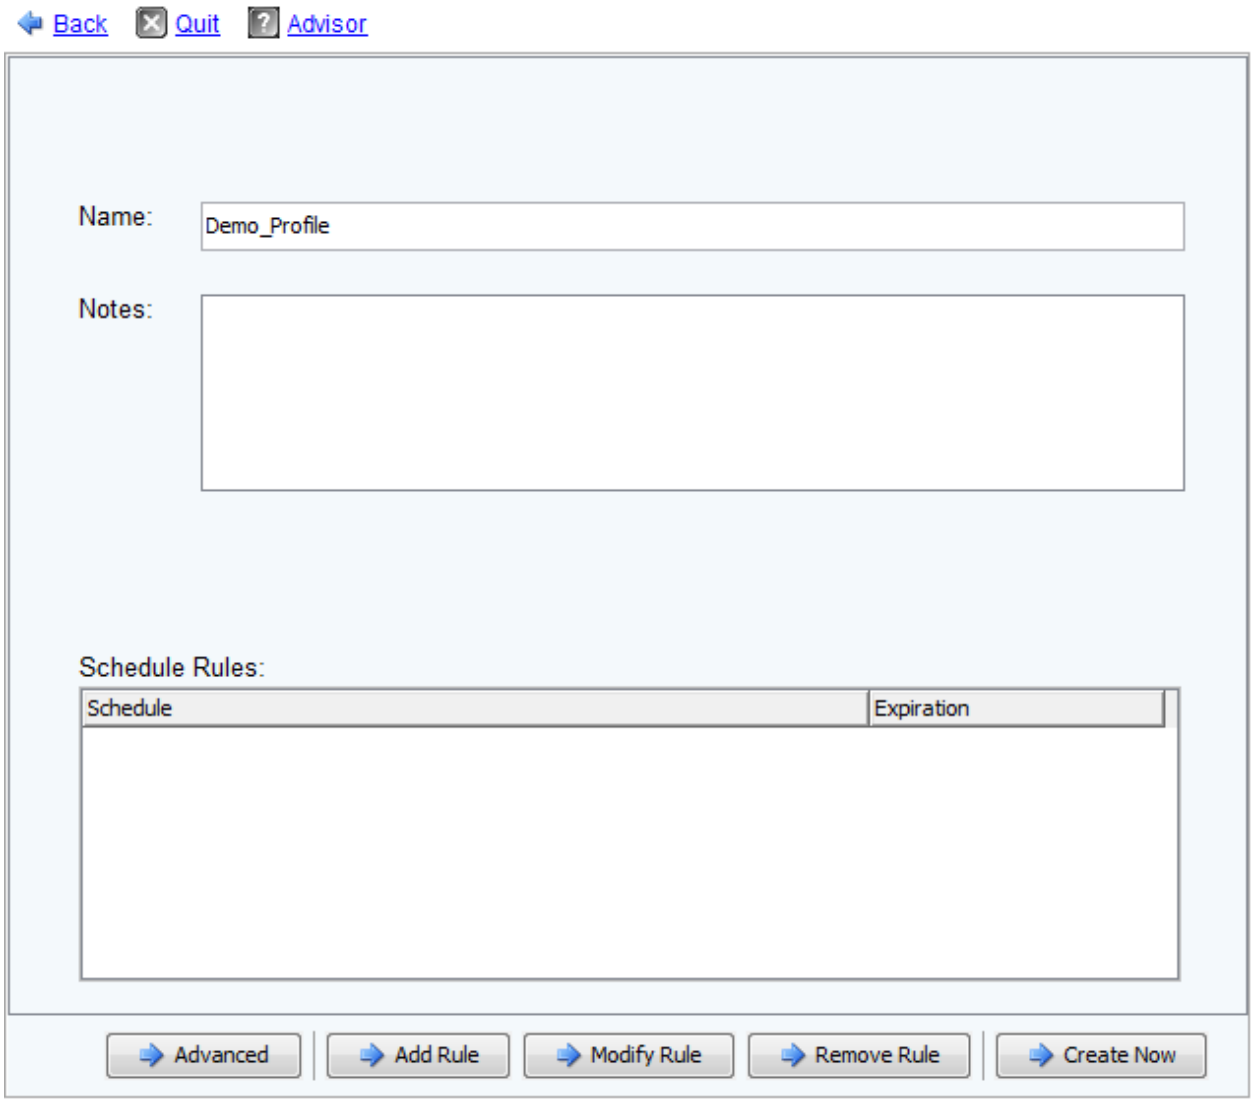

In this next dialog, left-Click the "Yes" button to create this Replay Profile without any scheduled rules.

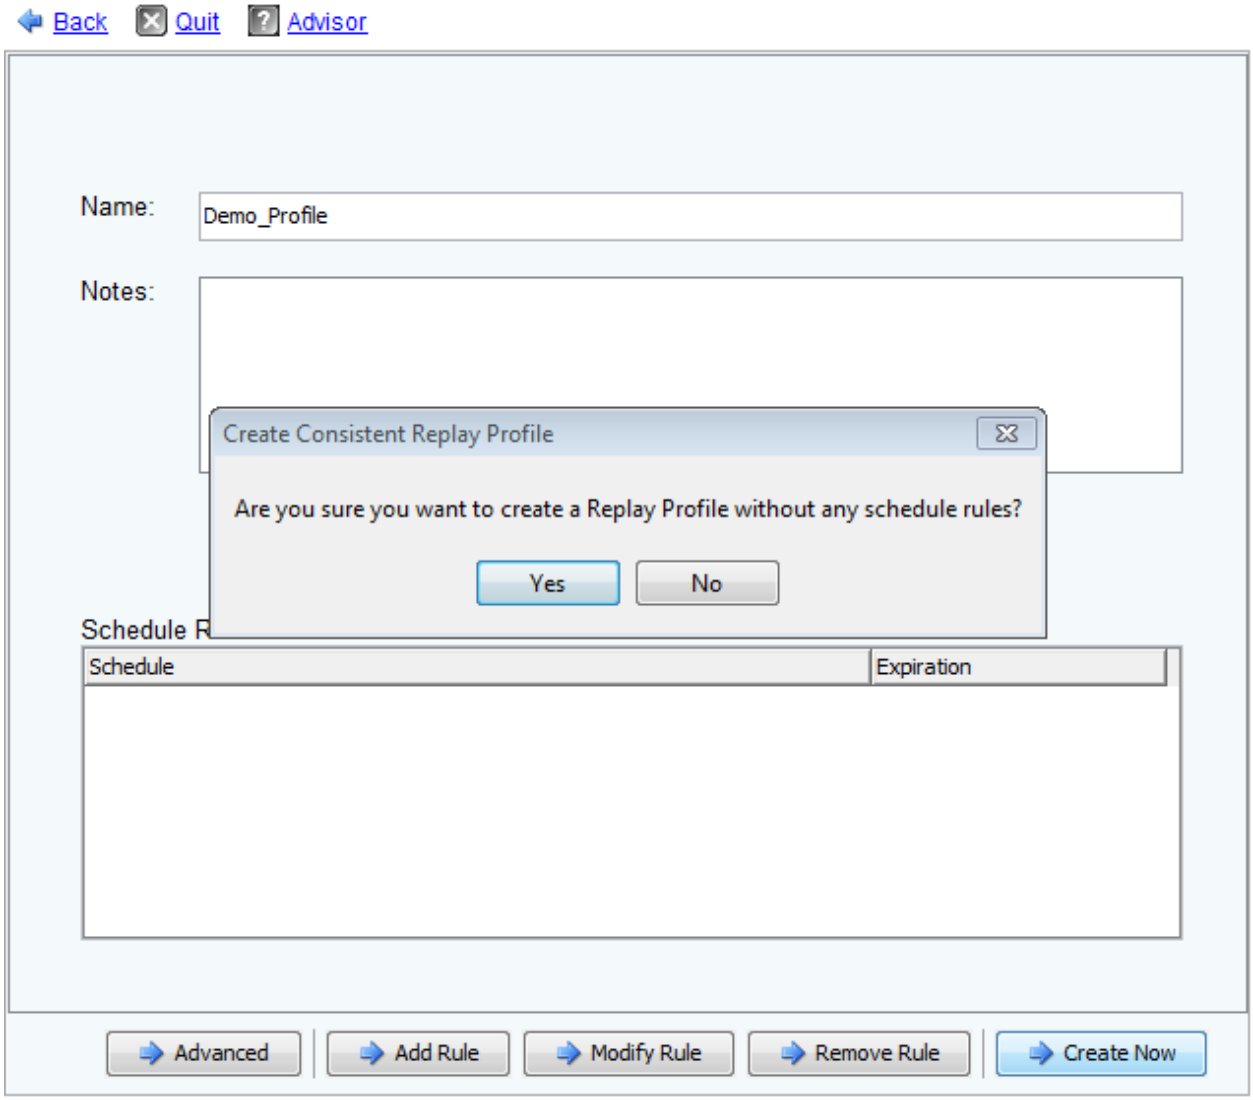

In this dialog, left-Click the “Apply to Volume(s)” button. This presents a subsequent dialog where you would assign this new Consistent Replay profile to all the volumes which are to be a member of this Consistency Group as seen below, then left-Click the “Continue” button.

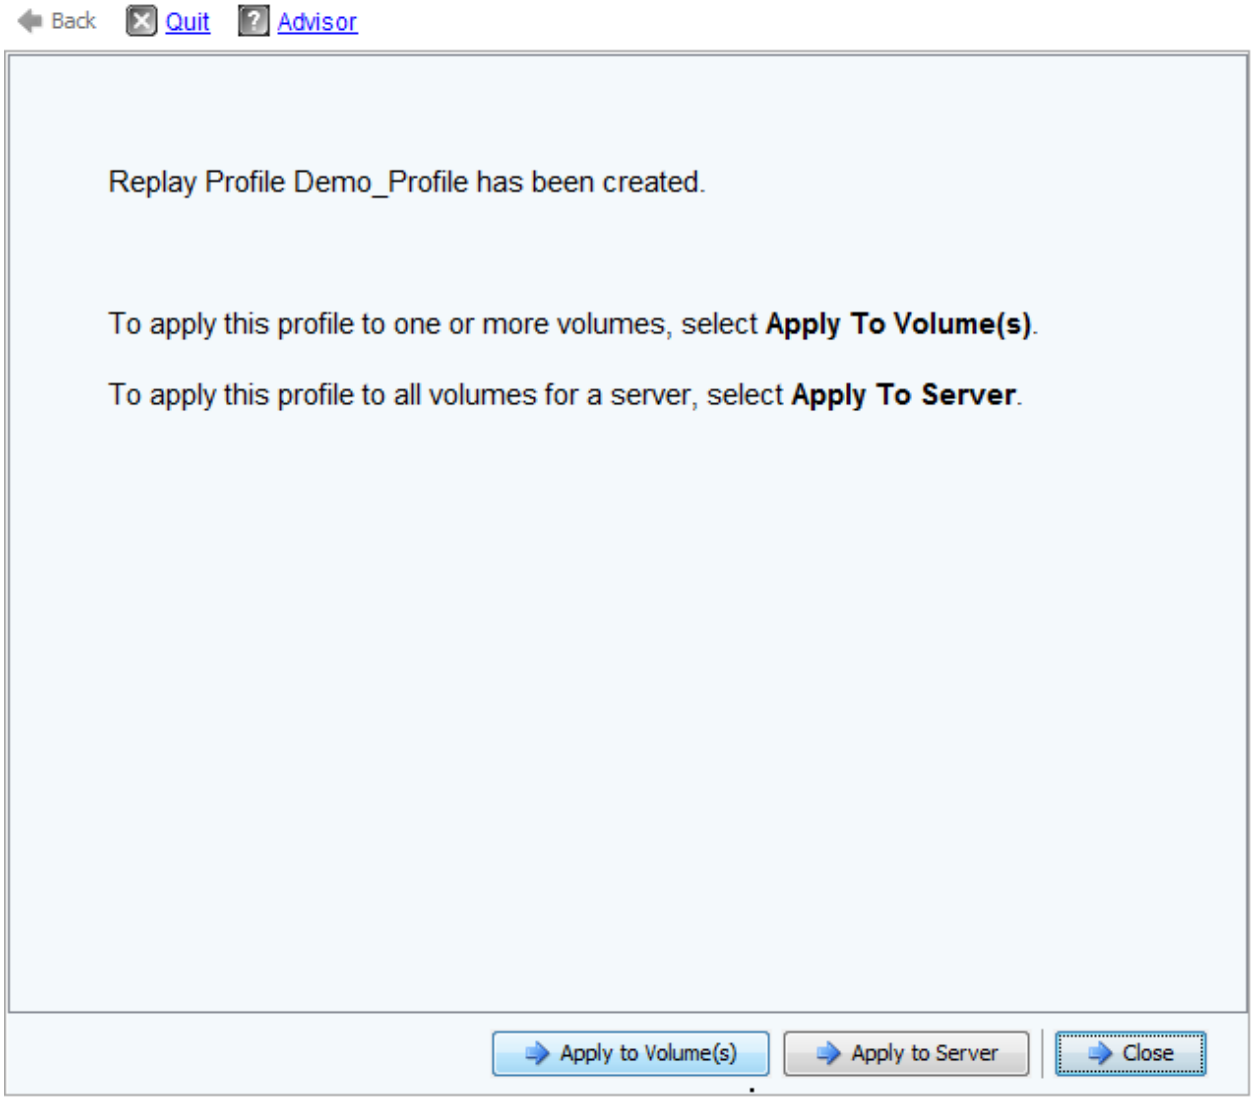
Back
Quit
Advisor
Replay Profile Demo_Profile has been created.
To apply this profile to one or more volumes, select Apply To Volume(s).
To apply this profile to all volumes for a server, select Apply To Server.
Apply to Volume(s)
Apply to Server
Close

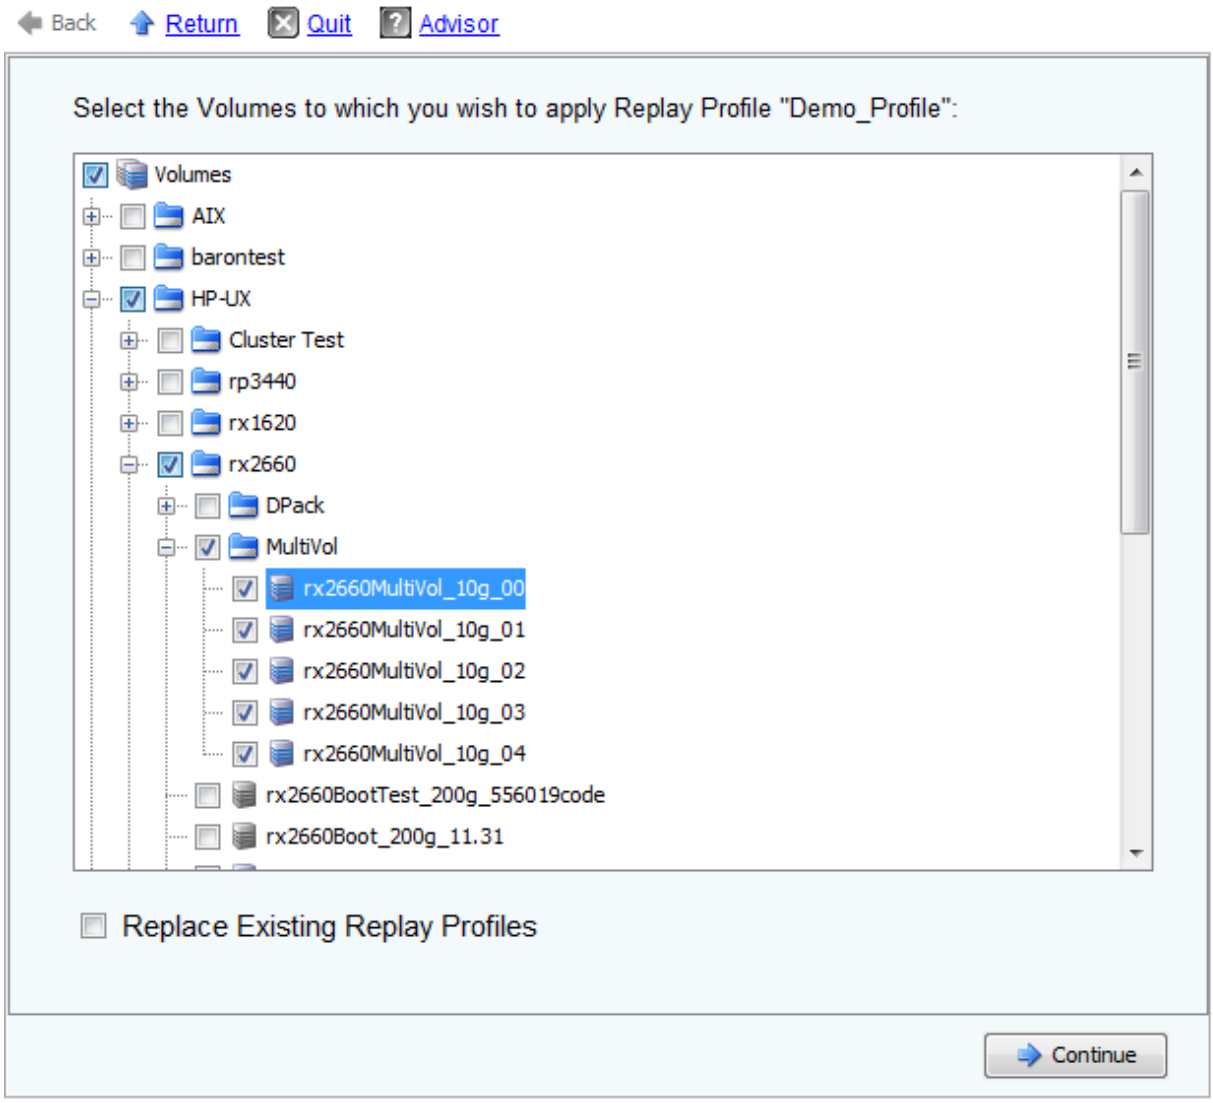

Finally in this last dialog, left-Click the “Apply Now” button to apply this Replay Profile to the selected Dell Compellent volumes.

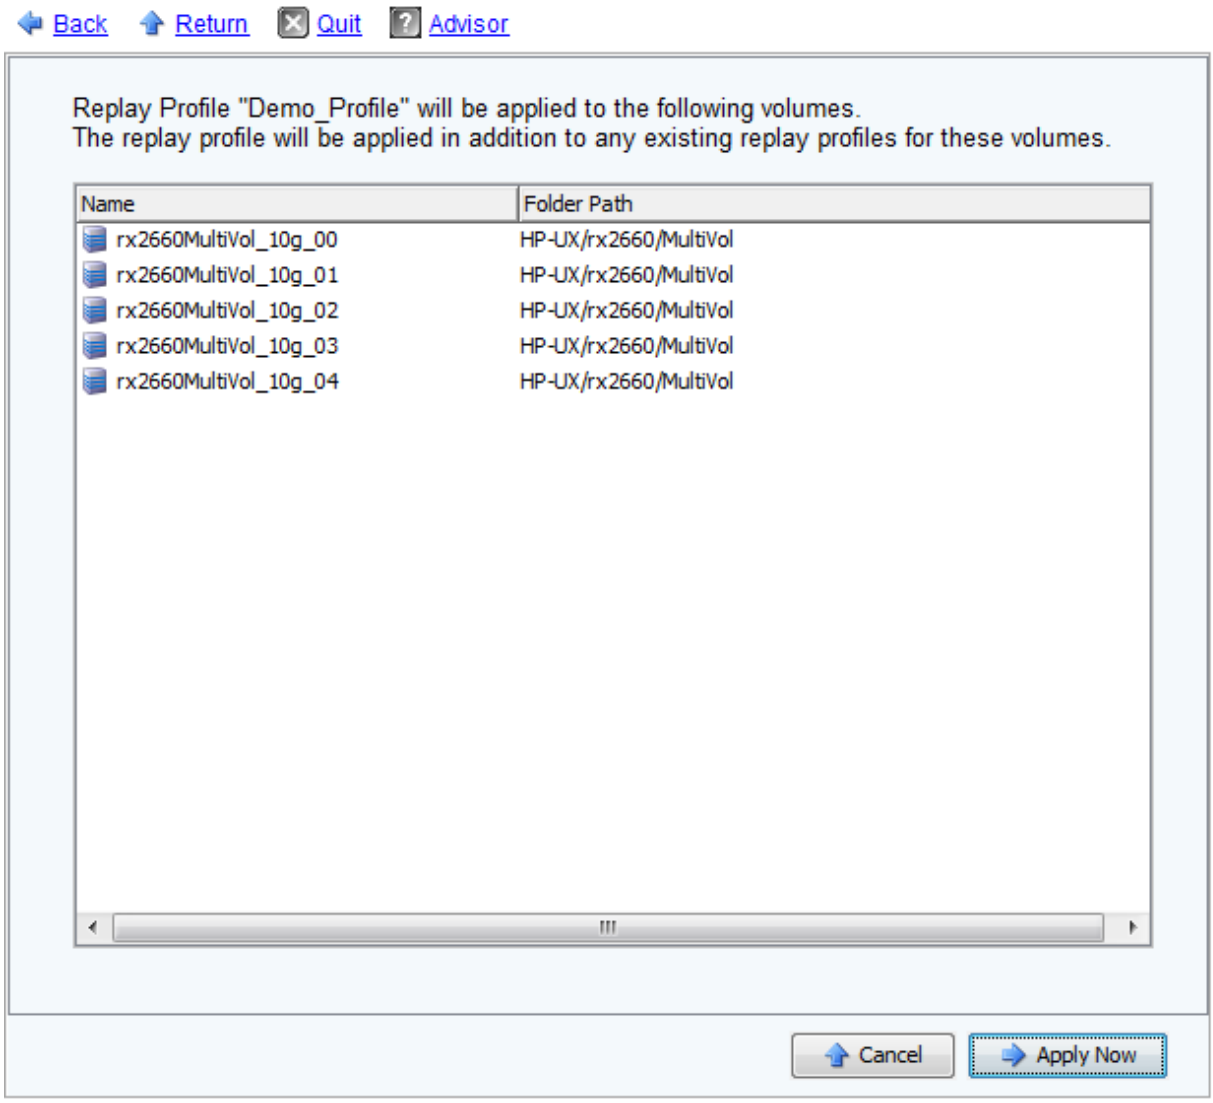

If importing to the same host, the vgchgid command needs to be applied to all volumes of the Volume Group at the same time.

Type: `vgchgid /dev/rdisk/disk147 /dev/rdisk/disk148 /dev/rdisk/disk149 /dev/rdisk/disk150 /dev/rdisk/disk156`

The vgimport command would likewise need to be applied to all volumes of the Volume Group at the same time as well.

Type: `vgimport –m /tmp/vgdemoStripe8_import.map –v vgdemoStripe8_import /dev/rdisk/disk147 /dev/rdisk/disk148 /dev/rdisk/disk149 /dev/rdisk/disk150 /dev/rdisk/disk156`

If you were mounting this same Volume Group back to the same host, you would not need to apply the mapfile as you see above (the mapfile is used when importing this Volume Group to a different target

host). The import of this Volume Group would then conclude with the appropriate vgchange, fsck, mkdir & mount commands accordingly.

Thin Provisioning

The concept of Thin Provisioning (and space reclamation) is the means by which an OS and its chosen volume and file management platform interact with the Dell Compellent Storage Center products. In this interaction, the OS (and volume/file management) communicates via SCSI library calls with the Dell Compellent Storage Center products and continually monitors the amount of space allocated (from the Dell Compellent Storage Center) versus the amount of space actually used. In doing so, the LUNs (as presented to the OS) can appear to be actually larger than the actual allocated (reserved) space on the Dell Compellent Storage Center. This creates a model of "just-enough" and "just-in-time" storage allocation, thereby resulting in huge savings both from a capital and operational perspective. For more information, please refer to this link http://h71028.www7.hp.com/enterprise/us/en/os/hpux11i-fsvm-overview.html.

HP-UX 11i v2 and older platform versions DOES NOT support Thin Provisioning interoperability with the Dell Compellent Storage Center products.

Tuning for Performance

In order to achieve optimized I/O with any OS platform, the proper tuning/configuration of various kernel, SCSI, network and filesystem parameters may be required; the process of tuning in itself is not a science but an art of understanding the performance and various I/O requirements of the respective environment and application ecosystem that you manage. Though the following information is shared as best practice for HP-UX with the Dell Compellent Storage Center products, it is neither a complete nor thorough representation of all possible tuning scenarios within any managed HP-UX environment. Consult with your respective vendor and software support engineers for your environment-centric tuning needs and for further assistance.

Kernel Tuning

Issue the following command against each the kernel, thereby increasing the SCSI Max Queue Depth to 32.

Type: `kctune -vs scsi_max_qdepth=32`

SCSI Configuration

Issue the following command against each Dell Compellent Storage Center-presented LUN, increasing its queue depth to 32.

Type: `scsictl –m queue_depth=32 /dev/rdsk/c27t0d0`

Filesystem Configuration

When creating a new filesystem on a Logical Volume, apply an 8k block size as part of the filesystem build options.

Type: `newfs –F vxfs –o largefiles –b 8192 /dev/vg<XX>/rlvol<name>`

Dynamic Root Disk

The HP-UX DRD (Dynamic Root Disk) toolset is made available by HP for purposes ranging from production/patch lifecycle management to OS upgrades and actual and/or exercised disaster recovery events. This toolset has been named (as such) for its intent to work primarily with your root Volume Group (i.e. vg00 under LVM control). This toolset is available for HP-UX 11i v3 (LVM 2.x and 1.x) and HP-UX 11i v2 (Sep 2004 Update release or newer, LVM 1.x); available for FREE download from https://h20392.www2.hp.com/portal/swdepot/displayProductInfo.do?productNumber=DynRootDisk.

This section will present the best use of Dell Compellent Storage Center presented LUNs with the DRD toolset. It is not meant to be a complete nor comprehensive discussion about the use of DRD for every HP-UX environment. For further assistance with DRD, reference the MAN pages included with the software package as well as the DRD Document Home located here http://h20000.www2.hp.com/bizsupport/TechSupport/DocumentIndex.jsp?lang=en&cc=us&taskId=101&prodClassId=10008&contentType=SupportManual&docIndexId=64255&prodTypeId=18964&prodSeriesId=4077178 and here http://bizsupport1.austin.hp.com/bc/docs/support/SupportManual/c01983454/c01983454.pdf.

With HP-UX 11i v2, you would install the following products to your target host and OS. This command may also need to reboot your system during the installation process in order to perform any necessary kernel modifications.

Type: `swinstall –x autoreboot=true –s /tmp/DRD_1123_WEB1201.depot \*`

(where the /tmp folder is the location where you had downloaded the .depot software package)

Upon installation, you would now have your DRD installation located here.

Type: `whereis drd`

```
drd: /opt/drd/bin/drd /opt/drd/share/man/man1m.Z/drd.1m
```

Type: `swlist –l bundle | grep –i dyn`

```
  DynRootDisk          B.1123.A.3.10.203 Dynamic Root Disk
```

The first step is to present a Dell Compellent Storage Center LUN to the target host (reference Server Setup, LUN Discovery). This LUN should be sized equivalent to or larger than the capacity of the root Volume Group. To discover the size of the vg00 Volume Group, run the command.

Type: `vgdisplay –v vg00 | grep PE`

```
Max PE per PV           4356
PE Size (Mbytes)        32
Total PE                4346
Alloc PE                1212
Free PE                 3134
```

(Where the multiplication of the two values in RED would give you the size of the Volume Group, in this case 32 * 4346 = 139072 ie. 146GB excl. overhead).

Having discovered the size of the vg00 Volume Group, we would now present a new LUN to the target host as shown below. For the sake of simplicity, we shall present a new 200GB LUN.

Type: `ioscan –fnC disk (2 items highlighted to reflect both paths to the same LUN)`

```
Class      I  H/W Path       Driver       S/W State   H/W Type     Description
===========================================================================
disk       5  0/0/2/1.0.16.0.0        sdisk        CLAIMED     DEVICE       TEAC    DV-28E-V
                      /dev/dsk/c0t0d0   /dev/rdsk/c0t0d0
disk       3  0/2/1/0/4/0.236.0.0.0.0.0 sdisk        CLAIMED     DEVICE       COMPELNTCompellent Vol
                      /dev/dsk/c1t0d0    /dev/dsk/c1t0d0s2   /dev/rdsk/c1t0d0    /dev/rdsk/c1t0d0s2
                      /dev/dsk/c1t0d0s1   /dev/dsk/c1t0d0s3   /dev/rdsk/c1t0d0s1
/dev/rdsk/c1t0d0s3
disk       8  0/2/1/0/4/0.236.0.0.0.0.1 sdisk        CLAIMED     DEVICE       COMPELNTCompellent Vol
                      /dev/dsk/c1t0d1   /dev/rdsk/c1t0d1
disk       6  0/2/1/0/4/0.236.1.0.0.0.1 sdisk        CLAIMED     DEVICE       COMPELNTCompellent Vol
                      /dev/dsk/c3t0d1   /dev/rdsk/c3t0d1
disk      11  0/2/1/0/4/0.236.1.0.0.0.2 sdisk        CLAIMED     DEVICE       COMPELNTCompellent Vol
                      /dev/dsk/c3t0d2   /dev/rdsk/c3t0d2
disk       4  0/2/1/0/4/1.236.7.0.0.0.0 sdisk        CLAIMED     DEVICE       COMPELNTCompellent Vol
                      /dev/dsk/c27t0d0    /dev/dsk/c27t0d0s2   /dev/rdsk/c27t0d0
/dev/rdsk/c27t0d0s2
                      /dev/dsk/c27t0d0s1   /dev/dsk/c27t0d0s3   /dev/rdsk/c27t0d0s1
/dev/rdsk/c27t0d0s3
disk       9  0/2/1/0/4/1.236.7.0.0.0.1 sdisk        CLAIMED     DEVICE       COMPELNTCompellent Vol
                      /dev/dsk/c27t0d1   /dev/rdsk/c27t0d1
disk       7  0/2/1/0/4/1.236.9.0.0.0.1 sdisk        CLAIMED     DEVICE       COMPELNTCompellent Vol
                      /dev/dsk/c31t0d1   /dev/rdsk/c31t0d1
disk      10  0/2/1/0/4/1.236.9.0.0.0.2 sdisk        CLAIMED     DEVICE       COMPELNTCompellent Vol
                      /dev/dsk/c31t0d2   /dev/rdsk/c31t0d2
disk       1  0/4/1/0.0.0.0.0         sdisk        CLAIMED     DEVICE       HP      DG146A4960
                      /dev/dsk/c41t0d0    /dev/dsk/c41t0d0s2   /dev/rdsk/c41t0d0
/dev/rdsk/c41t0d0s2
                      /dev/dsk/c41t0d0s1   /dev/dsk/c41t0d0s3   /dev/rdsk/c41t0d0s1
/dev/rdsk/c41t0d0s3
disk       2  0/4/1/0.0.0.1.0         sdisk        CLAIMED     DEVICE       HP      DG146BAAJB
                      /dev/dsk/c41t1d0    /dev/dsk/c41t1d0s2   /dev/rdsk/c41t1d0
/dev/rdsk/c41t1d0s2
                      /dev/dsk/c41t1d0s1   /dev/dsk/c41t1d0s3   /dev/rdsk/c41t1d0s1
/dev/rdsk/c41t1d0s3
```

DRD and creating a Point in Time Copy

The DRD toolset is a fantastic means by which to create a Point in Time Copy (henceforth referred to as "**PITC**") of the vg00 Volume Group. This PITC would most ideally be used as either a running

configuration snapshot of the system or as a backup copy of your vg00 Volume Group (where you might be installing a yet untested software package or patch set). Once created, this PITC may be brought back online as a mounted filesystem in cases where you might need to rollback certain files to previous known good versions.

To create a PITC, issue the command.

```
Type: /opt/drd/bin/drd clone –v –x overwrite=true –t /dev/dsk/c3t0d2

======= 11/21/11 11:31:42 CST BEGIN Clone System Image (user=root) (jobid=rx3600)

    * Reading Current System Information
    * Selecting System Image To Clone
    * Selecting Target Disk
    * Selecting Volume Manager For New System Image
    * Analyzing For System Image Cloning
    * Creating New File Systems
    * Copying File Systems To New System Image
    * Making New System Image Bootable
    * Unmounting New System Image Clone
    * System image: "sysimage_001" on disk "/dev/dsk/c3t0d2"

======= 11/21/11 11:55:42 CST END Clone System Image succeeded. (user=root) (jobid=rx3600)
```

Note

On an Itanium (ia64) based OS installations, this “drd clone” command will only replicate the s1 and s2 partitions to the cloned LUN. The s3 (HP Service Partition) is not replicated.

To mount the PITC, issue the command.

```
Type: /opt/drd/bin/drd mount (likewise, “drd umount” to umount the DRD clone
image)

Filesystem          kbytes    used   avail %used Mounted on
/dev/vg00/lvol3    1048576  322640  720320   31% /
/dev/vg00/lvol1    2097152   85800 1995744    4% /stand
/dev/vg00/lvol8    6029312 2218424 3783008   37% /var
/dev/vg00/lvol7    6914048 2663360 4217488   39% /usr
/dev/vg00/lvol6    4194304   16824 4144848    0% /tmp
/dev/vg00/lvol5   10518528 4065208 6402936   39% /opt
/dev/vg00/lvol4     524288   16696  503632    3% /home
/dev/drd00/lvol3   1048576  339640  703456   33% /var/opt/drd/mnts/sysimage_001
/dev/drd00/lvol4    524288   16696  503632    3% /var/opt/drd/mnts/sysimage_001/home
/dev/drd00/lvol5  10551296 4060264 6440352   39% /var/opt/drd/mnts/sysimage_001/opt
/dev/drd00/lvol1   2097152   85752 1995728    4% /var/opt/drd/mnts/sysimage_001/stand
/dev/drd00/lvol6   4194304   16816 4144856    0% /var/opt/drd/mnts/sysimage_001/tmp
/dev/drd00/lvol7   6946816 2657912 4255400   38% /var/opt/drd/mnts/sysimage_001/usr
```

```
/dev/drd00/lvol8   6029312 2200760 3798680   37% /var/opt/drd/mnts/sysimage_001/var
```

DRD and creating a Boot from SAN LUN

On the Dell Compellent Storage Center, locate the LUN in the "Storage -> Volumes" tree. We will right-Click on this LUN and choose the "Map Volume to Server" option. This will present the dialog as shown below (in this dialog, you would select the host that you wish to map this LUN to then left-Click the "Continue" button).

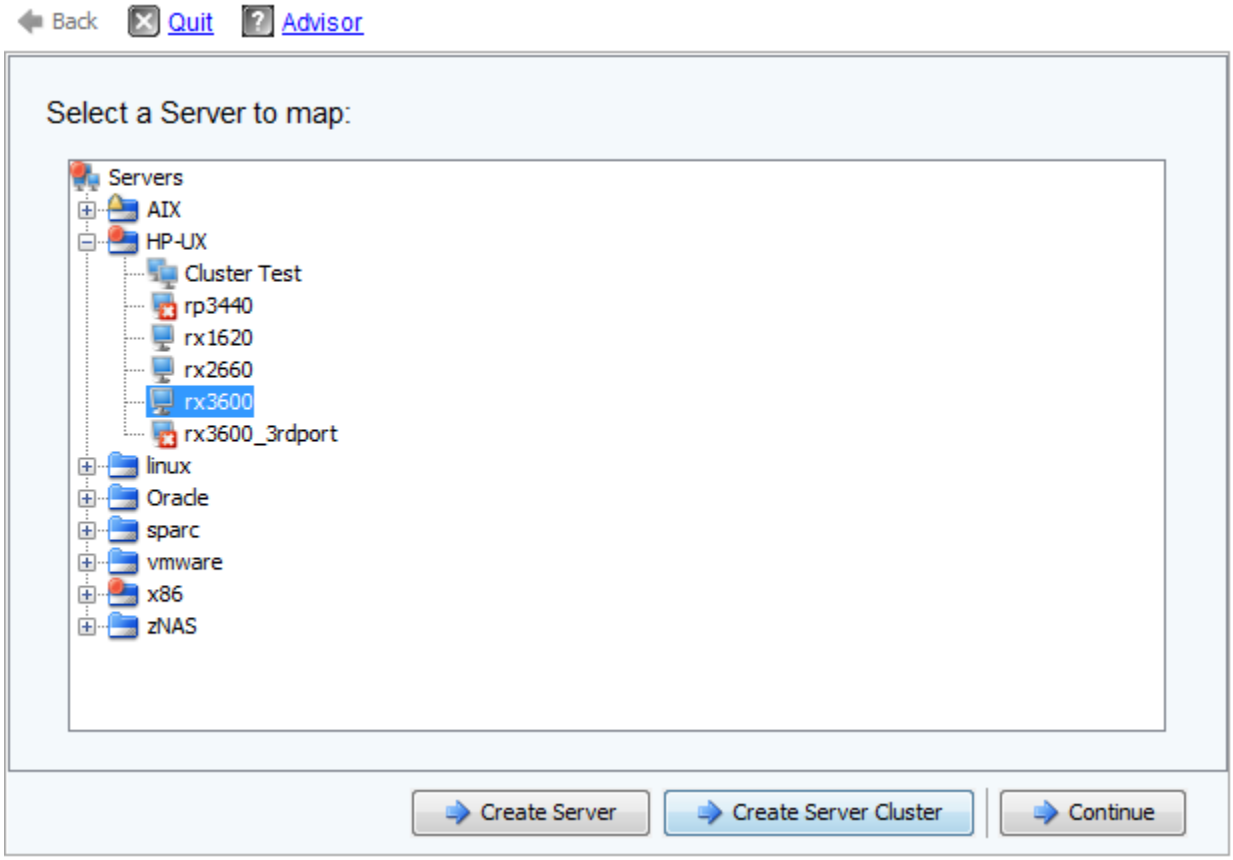

In this next dialog, you would left-Click on the "Advanced" button.

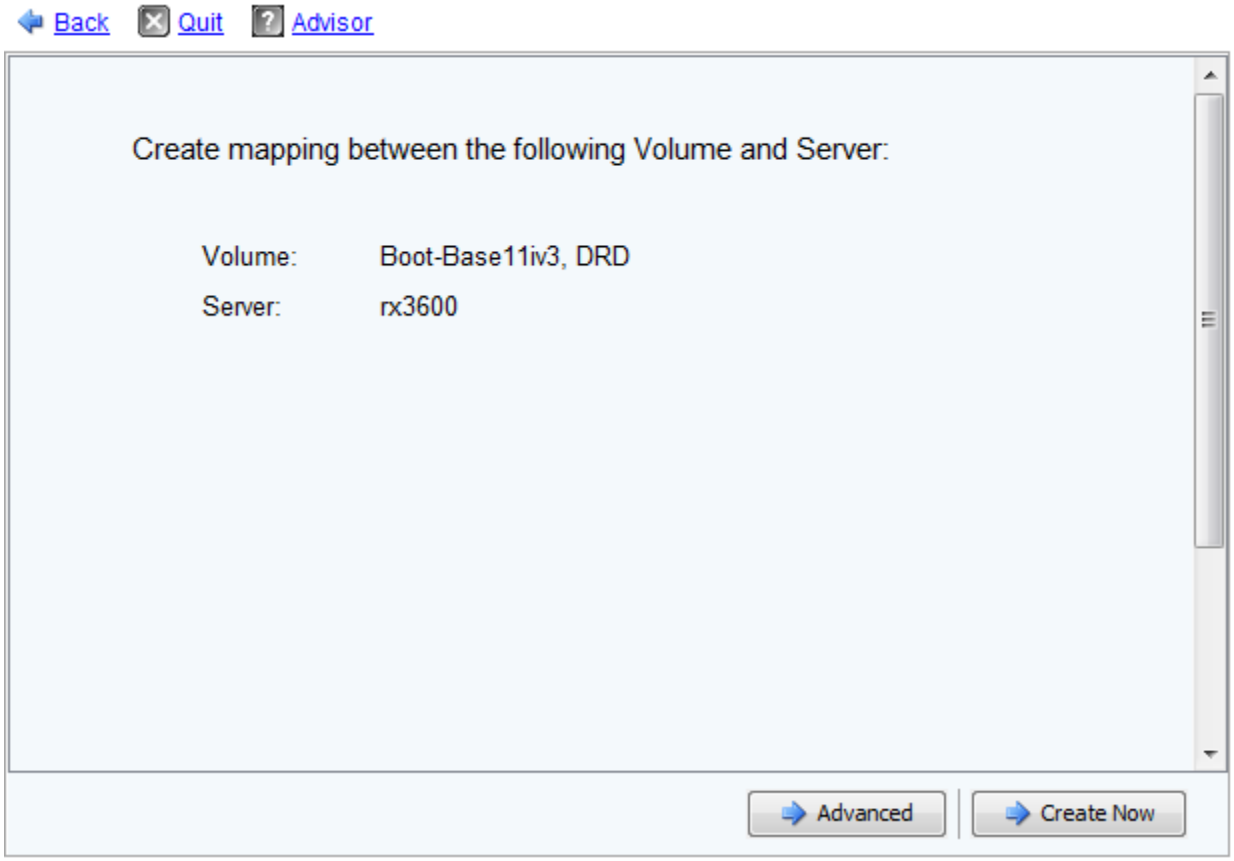

In this Advanced configuration dialog, you want to mark the checkbox which says “Create maps to down server ports”, this checkbox is only available if the FC ports to the target host is in a Down state; then left-Click the “Continue” button.

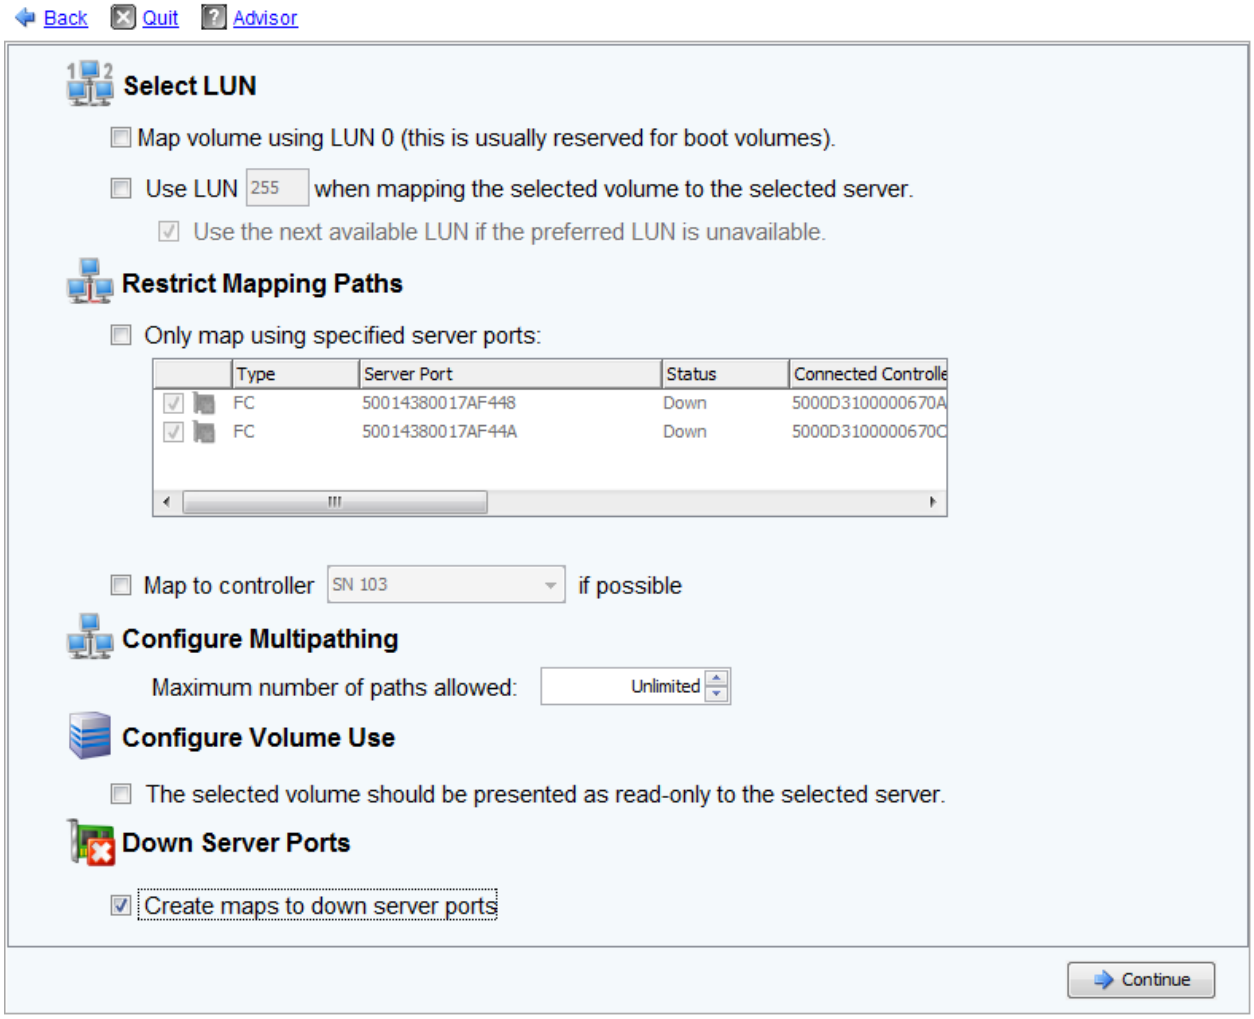

Finally, you would left-Click the "Create Now" button to map this LUN to the host.

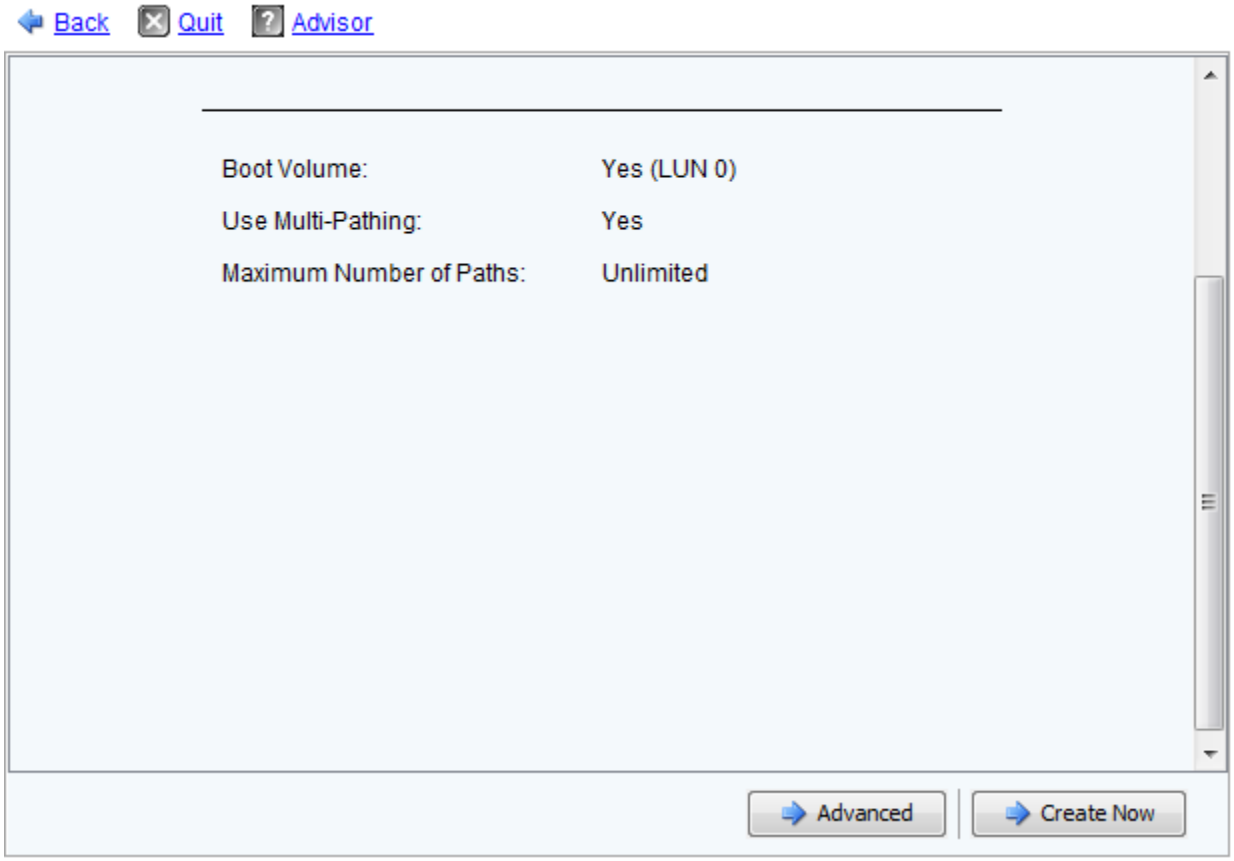

With this new bootable LUN mapped, you would now proceed to identifying this new LUN on the host and thereafter creating the drd clone as detailed in the section above. As shown in RED below (2 items highlighted to reflect both paths to the same LUN). Note: your LUN discovery, controller, target and lun paths may vary from the captured output shown below.

Type: `ioscan -fnC disk`

```
Class      I  H/W Path       Driver       S/W State   H/W Type    Description
==========================================================================
disk       5  0/0/2/1.0.16.0.0        sdisk        CLAIMED     DEVICE      TEAC    DV-28E-V
                     /dev/dsk/c0t0d0   /dev/rdsk/c0t0d0
disk       3  0/2/1/0/4/0.236.0.0.0.0.0  sdisk        CLAIMED     DEVICE      COMPELNTCompellent Vol
                     /dev/dsk/c1t0d0   /dev/dsk/c1t0d0s2   /dev/rdsk/c1t0d0   /dev/rdsk/c1t0d0s2
                     /dev/dsk/c1t0d0s1   /dev/dsk/c1t0d0s3   /dev/rdsk/c1t0d0s1
/dev/rdsk/c1t0d0s3
disk       8  0/2/1/0/4/0.236.0.0.0.0.1  sdisk        CLAIMED     DEVICE      COMPELNTCompellent Vol
                     /dev/dsk/c1t0d1   /dev/rdsk/c1t0d1
disk       6  0/2/1/0/4/0.236.1.0.0.0.1  sdisk        CLAIMED     DEVICE      COMPELNTCompellent Vol
                     /dev/dsk/c3t0d1   /dev/rdsk/c3t0d1
disk      11  0/2/1/0/4/0.236.1.0.0.0.2  sdisk        CLAIMED     DEVICE      COMPELNTCompellent Vol
                     /dev/dsk/c3t0d2   /dev/dsk/c3t0d2s2   /dev/rdsk/c3t0d2   /dev/rdsk/c3t0d2s2
```

```
                    /dev/dsk/c3t0d2s1   /dev/dsk/c3t0d2s3   /dev/rdsk/c3t0d2s1
/dev/rdsk/c3t0d2s3
disk        4  0/2/1/0/4/1.236.7.0.0.0.0  sdisk        CLAIMED     DEVICE      COMPELNTCompellent Vol
                    /dev/dsk/c27t0d0    /dev/dsk/c27t0d0s2   /dev/rdsk/c27t0d0
/dev/rdsk/c27t0d0s2
                    /dev/dsk/c27t0d0s1   /dev/dsk/c27t0d0s3   /dev/rdsk/c27t0d0s1
/dev/rdsk/c27t0d0s3
disk        9  0/2/1/0/4/1.236.7.0.0.0.1  sdisk        CLAIMED     DEVICE      COMPELNTCompellent Vol
                    /dev/dsk/c27t0d1   /dev/rdsk/c27t0d1
disk        7  0/2/1/0/4/1.236.9.0.0.0.1  sdisk        CLAIMED     DEVICE      COMPELNTCompellent Vol
                    /dev/dsk/c31t0d1   /dev/rdsk/c31t0d1
disk       10  0/2/1/0/4/1.236.9.0.0.0.2  sdisk        CLAIMED     DEVICE      COMPELNTCompellent Vol
                    /dev/dsk/c31t0d2    /dev/dsk/c31t0d2s2   /dev/rdsk/c31t0d2
/dev/rdsk/c31t0d2s2
disk        1  0/4/1/0.0.0.0.0          sdisk        CLAIMED     DEVICE      HP      DG146A4960
                    /dev/dsk/c41t0d0    /dev/dsk/c41t0d0s2   /dev/rdsk/c41t0d0
/dev/rdsk/c41t0d0s2
                    /dev/dsk/c41t0d0s1   /dev/dsk/c41t0d0s3   /dev/rdsk/c41t0d0s1
/dev/rdsk/c41t0d0s3
disk        2  0/4/1/0.0.0.1.0          sdisk        CLAIMED     DEVICE      HP      DG146BAAJB
                    /dev/dsk/c41t1d0    /dev/dsk/c41t1d0s2   /dev/rdsk/c41t1d0
/dev/rdsk/c41t1d0s2
                    /dev/dsk/c41t1d0s1   /dev/dsk/c41t1d0s3   /dev/rdsk/c41t1d0s1
/dev/rdsk/c41t1d0s3
```

The use of DRD with Veritas VxVM managed root filesystems (ie. rootdg) is outside the scope of this document; kindly refer to the HP Support Center for further instructions and/or assistance. Additional information may be found from this link below.
http://bizsupport1.austin.hp.com/bizsupport/TechSupport/Document.jsp?lang=en&cc=us&taskId=110&prodSeriesId=3203310&prodTypeId=18964&objectID=c01744440&printver=true

Boot from SAN (OS Boot Volume)

To boot from SAN, you would need to configure the Dell Compellent Storage Center presented LUN as a bootable device path. We would achieve this with the command as follows.

Type: `/opt/drd/bin/drd activate`

OR

Type: `setboot –p 0/2/1/0/4/0.236.0.0.0.0.0`

(Where **-p**: represents configuring the path as the primary boot device and where either of the dual paths may be used to identify the bootable device).

Upon reboot, the host will attempt to boot from this SAN LUN.

The HP-UX 11i v2 IgniteUX C.7.13.259 IINSTALLFS kernel (or newer) **CONTAINS** the necessary drivers to detect Dell Compellent Storage Center presented FC LUNs during a IgniteUX-based installation process. You may either install HP-UX directly to the Dell Compellent Storage Center LUN via the IgniteUX interactive GUI or the drd process (outlined in DRD and creating a Point in Time Copy, DRD and creating a Boot from SAN LUN).

HP-UX 11i v1 only operates on PA-RISC based architecture. This OS version **DOES NOT** provide any support for drd tools nor provide any facilities to boot from SAN LUNs.

iSCSI and HP-UX

iSCSI is an Internet Protocol (IP)-based storage networking standard for linking data storage facilities. By carrying SCSI commands over IP networks, iSCSI is used to facilitate data transfers over Intranets and to manage storage over long distances. iSCSI can be used to transmit data over local area networks (LANs), wide area networks (WANs), or the Internet and can enable location-independent data storage and retrieval.

Assumptions

- iSCSI ports are setup and configured on the Dell Compellent Storage Center
- iSCSI LAN HBAs have an assigned IP address
- The HP-UX iSCSI initiator software is downloaded and installed for your specific version of HP-UX
- Any relevant HP-UX iSCSI patches are applied

Setup

In order to validate that you have the proper software installed, run the command.

```
Type: swlist –l bundle | grep iSCSI

  iSCSI-00          B.11.23.03f    HP-UX iSCSI Software Initiator
```

Having validated the software installation, we now validate that the necessary processes (required for iSCSI to function, shown in RED) are running as well. Issue the command.

```
Type: ps –ef | grep iscsi

   root  1387    1  0 12:32:26 ?       0:00 /opt/iscsi/bin/iswd
   root  1385    1  0 12:32:26 ?       0:00 /opt/iscsi/bin/iscsi_resolvd
   root  2706    1  0 12:36:19 ?       0:00 /usr/sbin/stm/uut/bin/tools/monitor/dm_iscsi_adapter
   root  3208  2178  0 12:37:11 console   0:00 grep iscsi
```

In order to setup an iSCSI connection to the Dell Compellent Storage Center, we need to capture the iSCSI Qualified Name (iqn) from our host. Issue the following command to discover the iqn that was system (automatically) generated.

```
Type: /opt/iscsi/bin/iscsiutil -l

Initiator Name          : iqn.1986-03.com.hp:rx3600.1ab0bf92-18a9-11dd-af0f-3b63380fc243
Initiator Alias         :

Authentication Method    :
CHAP Method             : CHAP_UNI
Initiator CHAP Name      :
CHAP Secret             :
NAS Hostname            :
NAS Secret              :
```

```
Radius Server Hostname    :
Header Digest             : None,CRC32C (default)
Data Digest               : None,CRC32C (default)
SLP Scope list for iSLPD  :
```

The system generated iqn is shown in RED above. While this automatically generated iqn is adequate to the task of setting up the iSCSI connection, it is recommended to change this iqn to a string, which is both more manageable and more meaningful. There are some generally accepted rules to iqn naming, the biggest one being that everything in front of the colon (:) should NOT be modified. The recommendation thus stands to change the random alpha numeric string AFTER the hostname to the MAC address of the network HBA used for the iSCSI connection. Obtain the MAC address with the following command.

Type: `lanscan`

```
Hardware Station        Crd Hdw   Net-Interface  NM  MAC       HP-DLPI DLPI
Path     Address        In# State NamePPA        ID  Type      Support Mjr#
0/2/1/0/6/0 0x001E0BFDA12E 0   UP    lan0 snap0     1   ETHER     Yes     119
0/2/1/0/6/1 0x001E0BFDA12F 1   UP    lan1 snap1     2   ETHER     Yes     119
0/4/2/0  0x001A4B0804B2 2   UP    lan2 snap2     3   ETHER     Yes     119
0/5/1/0/6/0 0x001708D1611C 4   UP    lan4 snap4     4   ETHER     Yes     119
0/4/2/1  0x001A4B0804B3 3   UP    lan3 snap3     5   ETHER     Yes     119
```

The MAC address of the active (for iSCSI use) lan interface (in this case lan2) is shown in RED above. The change the iqn, issue the command.

Type: `/opt/iscsi/bin/iscsiutil –i –N iqn.1986-03.com.hp:rx3600.001A4B0804B2`

```
iscsiutil: Initiator Name "iqn.1986-03.com.hp:rx3600.001A4B0804B2" has been successfully updated.
```

To validate the change, we issue this command again.

Type: `/opt/iscsi/bin/iscsiutil –l | grep Initiator`

```
Initiator Name          : iqn.1986-03.com.hp:rx3600.001A4B0804B2
Initiator Alias         :
Initiator CHAP Name     :
```

We would now need to add the discovery IP address to the iSCSI configuration. The discovery IP address is the IP address assigned to the iSCSI I/O card (or Controller Ports) on your Dell Compellent Storage Center. For purposes of this example, we have one (1) VLAN setup (VLAN 550), this VLAN is mapped to a single Controller Port (akin to a Virtual IP interface) on the Dell Compellent Storage Center, where this Controller Port further represents four (4) configured iSCSI interfaces; two (2) of these interfaces on each controller head of your Dell Compellent Storage Center for redundancy. The discovery IP addresses for this Controller Port is 172.16.26.180. We now issue the command.

```
Type: /opt/iscsi/bin/iscsiutil –a –I 172.16.26.180

Target address "172.16.26.180:3260,1" has been successfully added.
```

If you have more than one (1) iSCSI IP address assigned to your Dell Compellent Storage Center, you would continue to add them in the same manner (ie. another Controller Port IP on another VLAN).

Additionally, it is noted that you should also ensure that the network interfaces on the host is appropriately connected to the proper Ethernet switches and zone these networks interfaces to the respective VLANs.

After you have added all of the discovery IP addresses, issue the following command to validate the discovery configuration.

```
Type: /opt/iscsi/bin/iscsiutil –p -D

Discovery Target Information
----------------------------

Target # 1
-----------
      IP Address              : 172.16.26.180
      iSCSI TCP Port          : 3260
      iSCSI Portal Group Tag    : 1

  User Configured:
  ----------------

      Authenticaton Method      :
      CHAP Method             : CHAP_UNI
      Initiator CHAP Name       :
      CHAP Secret             :
      Header Digest           : None,CRC32C (default)
      Data Digest             : None,CRC32C (default)
```

Next we issue the following command, to register this new iqn to the Dell Compellent Storage Center.

```
Type: ioscan –fnH 255/0

Class     I  H/W Path   Driver   S/W State   H/W Type     Description
=============================================================================
iscsi     0  255/0      iscsi    CLAIMED     VIRTBUS      iSCSI Virtual Node
ext_bus  42  255/0/0.0  iscsial  CLAIMED     INTERFACE    iSCSI-SCSI Protocol Interface
ext_bus  43  255/0/1.0  iscsial  CLAIMED     INTERFACE    iSCSI-SCSI Protocol Interface
ext_bus  44  255/0/2.0  iscsial  CLAIMED     INTERFACE    iSCSI-SCSI Protocol Interface
```

```
ext_bus  45  255/0/3.0  iscsial  CLAIMED     INTERFACE    iSCSI-SCSI Protocol Interface
```

In the scenario above, we have only setup only one (1) discovery IP address mapped to a single interface on a quad port NIC. To ensure proper failover and redundancy in an enterprise scenario, it is strongly recommended to establish one or more discovery IP addresses on the host (OS) side; in accordance with mapping a secondary interface on a alternate NIC via a different VLAN to the Dell Compellent Storage Center.

Adding LUNs

Having completed the procedures above, you can now add this connection to the Dell Compellent Storage Center. Open the "Create Server" dialog and uncheck the "Only Show Up Connections" box and let the screen refresh. The iqn for your server should appear in the dialog box.

If the iqn does not appear this latter dialog box, you will need to add the iqn manually. Left-click the "Manually Define HBA" button, select the iSCSI radio button in the "Transport Type" selection box and enter the known iqn into the "WWN or iSCSI Name:" text box, left-Click the "Continue" button, then continue creating the server object per usual procedures. This procedure should NOT be needed if using Dell Compellent Storage Center code 5.4.x or newer.

Complete the "Create Server" process, then map LUNs to the host. For purposes of this example, we have created and mapped two (2) 50GB LUNs to the host. On the host, issue this command.

Type: `ioscan –fnH 255/0`

```
Class     I  H/W Path      Driver   S/W State   H/W Type     Description
=============================================================================
iscsi     0  255/0         iscsi    CLAIMED     VIRTBUS      iSCSI Virtual Node
ext_bus  42  255/0/0.0     iscsial  CLAIMED     INTERFACE    iSCSI-SCSI Protocol Interface
target   44  255/0/0.0.0   tgt      CLAIMED     DEVICE
disk      3  255/0/0.0.0.1 sdisk    CLAIMED     DEVICE       COMPELNTCompellent Vol
ext_bus  43  255/0/1.0     iscsial  CLAIMED     INTERFACE    iSCSI-SCSI Protocol Interface
target   45  255/0/1.0.0   tgt      CLAIMED     DEVICE
disk      4  255/0/1.0.0.1 sdisk    CLAIMED     DEVICE       COMPELNTCompellent Vol
ext_bus  44  255/0/2.0     iscsial  CLAIMED     INTERFACE    iSCSI-SCSI Protocol Interface
target   46  255/0/2.0.0   tgt      CLAIMED     DEVICE
disk     12  255/0/2.0.0.1 sdisk    CLAIMED     DEVICE       COMPELNTCompellent Vol
ext_bus  45  255/0/3.0     iscsial  CLAIMED     INTERFACE    iSCSI-SCSI Protocol Interface
target   47  255/0/3.0.0   tgt      CLAIMED     DEVICE
disk     13  255/0/3.0.0.1 sdisk    CLAIMED     DEVICE       COMPELNTCompellent Vol
```

We see the LUNs presented from the Dell Compellent Storage Center to the host and would now need to build the necessary devices files to support these LUNs. Issue the command.

Type: `insf -e`

Finally, we issue this command and observe that all LUNs have their corresponding devices files in place.

Type: `ioscan –fnH 255/0`

```
Class     I  H/W Path      Driver   S/W State   H/W Type     Description
===================================================================
iscsi     0  255/0         iscsi    CLAIMED     VIRTBUS      iSCSI Virtual Node
ext_bus  42  255/0/0.0     iscsial  CLAIMED     INTERFACE    iSCSI-SCSI Protocol Interface
target   44  255/0/0.0.0   tgt      CLAIMED     DEVICE
disk      3  255/0/0.0.0.1 sdisk    CLAIMED     DEVICE       COMPELNTCompellent Vol
                  /dev/dsk/c42t0d1   /dev/rdsk/c42t0d1
ext_bus  43  255/0/1.0     iscsial  CLAIMED     INTERFACE    iSCSI-SCSI Protocol Interface
target   45  255/0/1.0.0   tgt      CLAIMED     DEVICE
disk      4  255/0/1.0.0.1 sdisk    CLAIMED     DEVICE       COMPELNTCompellent Vol
                  /dev/dsk/c43t0d1   /dev/rdsk/c43t0d1
ext_bus  44  255/0/2.0     iscsial  CLAIMED     INTERFACE    iSCSI-SCSI Protocol Interface
target   46  255/0/2.0.0   tgt      CLAIMED     DEVICE
disk     12  255/0/2.0.0.1 sdisk    CLAIMED     DEVICE       COMPELNTCompellent Vol
                  /dev/dsk/c44t0d1   /dev/rdsk/c44t0d1
ext_bus  45  255/0/3.0     iscsial  CLAIMED     INTERFACE    iSCSI-SCSI Protocol Interface
target   47  255/0/3.0.0   tgt      CLAIMED     DEVICE
disk     13  255/0/3.0.0.1 sdisk    CLAIMED     DEVICE       COMPELNTCompellent Vol
                  /dev/dsk/c45t0d1   /dev/rdsk/c45t0d1
```

The devices are now ready for use (apply normal LVM procedures to create required Volume Group, Logical Volumes and filesystems therefrom). For purposes of this example, these four (4) LUNs were added to a new Volume Group vg04 whereby the results look as follow.

Type: `vgdisplay –v vg04`

```
<snip>

   --- Physical volumes ---
   PV Name                  /dev/dsk/c42t0d1
   PV Name                  /dev/dsk/c43t0d1 Alternate Link
   PV Status                available
   Total PE                 12797
   Free PE                  12797
   Autoswitch               On
   Proactive Polling        On

   PV Name                  /dev/dsk/c44t0d1
   PV Name                  /dev/dsk/c45t0d1 Alternate Link
   PV Status                available
```

Total PE	12797
Free PE	12797
Autoswitch	On
Proactive Polling	On

CompCU (Dell Compellent Command Utility)

Dell Compellent has available (for download) a java-based command utility (CompCU) that can perform most of the functions that can be performed from the GUI. This tool allows for the scripting and automated integration of SAN tasks between the HP-UX operating system and the Dell Compellent Storage Center. CompCU can be used to script common administrative tasks that can tremendous time savers and provide a consistent framework for managing volumes and replays. If you decide to obtain and use this utility, this section should help. This section below is not intended to be a comprehensive representation of all CompCU functions, but meant as an initial demonstration and offer insight into the sort of tasks that CompCU can accomplish.

Setup and Validation

1. Obtain the CompCU utility from the Dell Compellent Knowledge Center Extranet Site http://kcint.compellent.com/Pages/home.aspx.

    a. Click: Software under the Downloads heading (in the left column).

    b. Select and download by clicking on the link "**Command Utility 5.4.1**" http://kcint.compellent.com/Pages/Download.aspx?DocID=1081.

    c. Unpack the zip file and transfer the "**CompCU.jar**" file to your HP-UX server.

    d. Familiarize yourself with the CompCU User Guide that comes in the zip file with the CompCU utility!

2. Ensure that you have a version of the Java JRE (formerly known as RTE) loaded on your HP-UX host which will work with your version of CompCU.

    a. Issue the following command.

    Type: `swlist –l bundle | grep –i java`

    ```
    Java15JDK        1.5.0.16.00   Java 1.5 JDK for HP-UX
    Java15JRE        1.5.0.16.00   Java 1.5 JRE for HP-UX
    Java60JDK        1.6.0.04.00   Java 6.0 JDK for HP-UX
    Java60JRE        1.6.0.04.00   Java 6.0 JRE for HP-UX
    T1456AA          1.4.2.22.00   Java2 1.4 SDK for HP-UX
    T1457AA          1.4.2.22.00   Java2 1.4 RTE for HP-UX
    T1458AA          1.4.2.22.00   Java2 1.4 Netscape Plugin for HP-UX
    ```

    Type: `whereis java`

    ```
    java: /opt/java1.5/bin/java /opt/java1.4/bin/java /opt/java6/bin/java
    /opt/java1.4/man/man1/java.1
    ```

    b. With Dell Compellent Storage Center 5.x, you would need java 1.6.x.x or greater to function with CompCU.

Multiple versions of java can be loaded on an HP-UX server without any issues. Any version of java on the system that you wish to use has to be referenced directly by its absolute path; there are no 'symbolic links' to from the java JRE binaries to "/usr/local/bin" nor is there any default PATH environment variables defined for java JRE binaries. The java installation is a **no reboot required** installation. The latest java versions (both JDK and JRE) for HP-UX may be obtained from http://software.hp.com.

3. At this point, if you have a good java JRE version installed and the CompCU utility on the system, you can issue a command against the Dell Compellent Storage Center. It is important to note that the server that you issue the CompCU commands from does not have to be connected to the Dell Compellent Storage Center nor have any LUNs mapped to it.

    a. Your command will look something like this:

    ```
    /opt/java6/bin/java -jar /tmp/CompCU.jar -host SC101 -user JoeAdmin -
    password PASSWORD -c "server show"
    ```

    b. The arguments are as follows:

        i. "-jar" needs the location of the CompCU.jar file.

        ii. "-host" needs a resolvable hostname for the target Dell Compellent Storage Center

        iii. "-user" needs a valid user name on the target Dell Compellent Storage Center; while you *can* use your Admin userid it is **very strongly discouraged!**

        iv. "-password" needs to be the password for the user name given for the "-user" argument.

        v. "-c" specifies the Dell Compellent Storage Center command to execute which follows in quotes.

    c. The output from the command shown in 3.a above should look something like this:

```
Compellent Command Utility (CompCU) 5.2.1.2

================================================================================
========
User Name:            JoeAdmin
Host/IP Address:      SC101
Single Command:       server show
================================================================================
========
Connecting to Storage Center: SC101 with user: JoeAdmin
Running Command: server show

Index  Name                    Folder                                    OsIndex  OS
Type              ParentIndex Parent                WWN_List
------ ------------------------------ ---------------------------------------------------------- --------
------ ---------------------- ---------- ------------------------------ ----------------------------------------
```

```
---------------------------------------------- 11    Apache                    Microsoft
2       Other Singlepath    Physical                                    210000E08B9D3B2C
12    Nova                  Microsoft                                   2       Other Singlepath
Physical                                    210000E08B9DAD2C
16    Impala-1              Netware                                     4       Other NoLUNGap
Physical                                    210000E08B89D7D6
20    Prowler               VMware/VMware_1                             33      VMWare ESX
3.5     Physical                                    210000E08B9D6B2C
93    Vantage               Jim's                                       2       Other Singlepath
Physical                                    2100001B32026455, 2101001B32226455
11      Compellent Any      Remote Storage Center                       5000D31000036001,
5000D31000036005, 5000D31000036009, 5000D3100003600D, 5000D31000036002, 199   vsan2 titan fc
interfaces      HP-UX/Titan                                 15      HP UX 11i v3        Physical
50060B00003082A6, 50060B0000226936                                  208
titan.0011855F0E6D          HP-UX/Titan                             15      HP UX 11i v3
Physical                                    iqn.1986-03.com.hp:itanic.0011855F0E6D
212   SQL2008-demo          Microsoft                                   9       Windows 2008
Physical                                    AC10169400000000
215   Tacoma                VMware/VMware_3                             34      VMWare ESX
4.0     Physical                                    210000E08B86FF58, 210000E08B862199

Successfully finished running Compellent Command Utility (CompCU) application.
```

4. Finally, it is once again and strongly recommended to familiarize yourself with the CompCU User Guide that comes in the zip file with the CompCU utility.

Example 1: Creating a Single Volume

The example below demonstrates using CompCU to create a volume on the Dell Compellent Storage Center. A 100GB volume named "100gTest" is created on the "sc10" Storage Center, in the "HP-UX" folder, and then mapped to the HP-UX host "rx2660".

```
# /opt/java6/bin/java -jar /tmp/CompCU.jar –host sc10.techsol.beer.town –user Admin –password
PASSWORD -c "volume create -name 100gTest -folder HP-UX -server rx2660 -size 100g"
Compellent Command Utility (CompCU) 5.5.1.4

================================================================================
User Name:            Admin
Host/IP Address:      sc10
Single Command:       volume create -name 100gTest -folder HP-UX -server rx2660 -size 100g
================================================================================
Connecting to Storage Center: sc10 with user: Admin
Running Command: volume create -name 100gTest -folder HP-UX -server rx2660 -size 100g
Creating Volume using StorageType 1: storagetype='Assigned-Redundant-4096', redundancy=Redundant,
pagesize=4096, diskfolder=Assigned.
Successfully mapped Volume '100gTest' to Server 'rx2660'
Successfully created Volume '100gTest', mapped it to Server 'rx2660' on Controller 'SN 103'
```

```
Successfully finished running Compellent Command Utility (CompCU) application.
```

Example 2: Creating Multiple Volumes

This next example demonstrates rapid deployment of several volumes and mapping them to the HP-UX host “rx2660”.

```
# for a in 0 1 2 3 4 5 6 7 8 9; do /opt/java6/bin/java -jar /tmp/CompCU.jar –host
sc10.techsol.beer.town –user Admin –password PASSWORD -c "volume create -folder HP-UX -lun 10${a} -
name 100gTest10${a} -server rx2660 -size 100g"; done
Compellent Command Utility (CompCU) 5.5.1.4

================================================================================
User Name:            Admin
Host/IP Address:      sc10
Single Command:       volume create -folder HP-UX -lun 100 -name 100gTest100 -server rx2660 -size
100g
================================================================================
Connecting to Storage Center: sc10 with user: Admin
Running Command: volume create -folder HP-UX -lun 100 -name 100gTest100 -server rx2660 -size 100g
Creating Volume using StorageType 1: storagetype='Assigned-Redundant-4096', redundancy=Redundant,
pagesize=4096, diskfolder=Assigned.
Successfully mapped Volume '100gTest100' to Server 'rx2660'
Successfully created Volume '100gTest100', mapped it to Server 'rx2660' on Controller 'SN 103'

Successfully finished running Compellent Command Utility (CompCU) application.

<snip>

Compellent Command Utility (CompCU) 5.5.1.4

================================================================================
User Name:            Admin
Host/IP Address:      sc10
Single Command:       volume create -folder HP-UX -lun 109 -name 100gTest109 -server rx2660 -size
100g
================================================================================
Connecting to Storage Center: sc10 with user: Admin
Running Command: volume create -folder HP-UX -lun 109 -name 100gTest109 -server rx2660 -size 100g
Creating Volume using StorageType 1: storagetype='Assigned-Redundant-4096', redundancy=Redundant,
pagesize=4096, diskfolder=Assigned.
Successfully mapped Volume '100gTest109' to Server 'rx2660'
Successfully created Volume '100gTest109', mapped it to Server 'rx2660' on Controller 'SN 104'

Successfully finished running Compellent Command Utility (CompCU) application.
```

Example 3: Capturing a Volume Replay

CompCU can also be used to create a Replay of the volume created above:

```
# /opt/java6/bin/java -jar /tmp/CompCU.jar –host sc10.techsol.beer.town –user Admin –password
PASSWORD -c "replay create -lun 200 -volume 100gTest100"
Compellent Command Utility (CompCU) 5.5.1.4

================================================================================
User Name:              Admin
Host/IP Address:        sc10
Single Command:         replay create -lun 200 -volume 100gTest100
================================================================================
Connecting to Storage Center: sc10 with user: Admin
Running Command: replay create -lun 200 -volume 100gTest100
Creating replay 'CUReplay_66022' on Volume '100gTest100' with no expiration

Successfully finished running Compellent Command Utility (CompCU) application.
```

FAQ